Climate Change and Field-Level Crop Quality, Yield, and Revenue

Sarah C. Smith* & Timothy K.M. Beatty†

December 12, 2024

Abstract

We quantify the effect of weather and climate on the revenue of processing-tomato farmers through yield and quality—quality being an understudied channel despite its role in price determination. Screening out low-quality products introduces selection bias into estimates of the effect of weather and climate on agriculture. Our novel data allow us to estimate this bias. We find that extreme temperatures reduce both yield and quality, leading to reduced revenue. While the yield effect dominates, failing to account for quality leads to a significant underestimate of the effect of temperature exposure on revenue. We predict climate change will significantly reduce yield and quality by century's end absent adaptation and all else equal. Yield effects are overstated while quality effects are understated when estimation relies on data on a subset of output that exceeds a quality threshold. Empirical work that ignores selection on quality may misrepresent the climate change challenge.

Keywords: quality, selection bias, irrigated agriculture, climate change

JEL Codes: Q12, Q54, C23

*Postdoctoral Scholar, Department of Agricultural and Resource Economics, University of California, Davis, scjsmith@ucdavis.edu.

†Professor, Department of Agricultural and Resource Economics, University of California, Davis, tbeatty@ucdavis.edu.

1 Introduction

The value of every agricultural product depends on its quality. Grain, meat, and milk are graded according to USDA quality standards, and fresh produce is sorted by size, color, and defects. Products grown under contract often face quality incentives, and low quality can violate contractual obligations or make products unmarketable. But work quantifying the economic impacts of extreme weather and climate change on agricultural production focuses almost exclusively on yield (see Carter et al., 2018 or Ortiz-Bobea, 2021 for summaries).

An important but understudied aspect of quality is its role in determining measured yield. For many agricultural products, data are only collected for the subset of output that is selected to be graded or harvested because it exceeds a minimum quality threshold. Selection on quality biases unconditional measures of both quality *and* yield, which limits a researcher's ability to recover the true effect of weather and climate change on agricultural outcomes. In short, ignoring quality biases estimates of the impact of weather and climate change on agricultural productivity and farm income.

We ask two related questions: How does weather affect the revenues of specialty-crop producers through both yield and quality, and what are the relative magnitudes of these effects? To what extent does ignoring selection on quality bias estimates of the effects of weather and climate change? We answer these questions using 13,000 field-level observations of processing-tomato yield, quality, and grower practices from across California between 2011 and 2021. These data are collected by a large tomato processor for the purposes of contracting and payment and capture the decisions of hundreds of commercial farmers. We use standard gridded weather data from PRISM, the well-established panel model (Deschênes & Greenstone, 2007; Schlenker & Roberts, 2009), and a restricted cubic-spline specification to flexibly estimate the effect of temperature exposure on yield, quality, and revenue.

This paper contributes to an emerging literature studying the effect of weather and climate change on agricultural product quality. Kawasaki & Uchida (2016), Kawasaki (2019), Dalhaus et al. (2020), and Ramsey et al. (2020) all find negative, economically important effects of weather on grower revenue with quality being a key pathway. This paper adds to this body of work by considering the pervasive issue of sample-selection bias introduced by the way yield and quality data are generated. To the best of our knowledge, only Kawasaki & Uchida (2016) control for the consequences of selection on quality in earlier work that estimated weather and climate change impacts on agricultural outcomes. Our setting and novel data offer advantages relative to earlier work.

First, our measures of quality have economic significance. Growers are paid a price per ton that depends on observed quality attributes, which introduces variation in price larger than price variation from market conditions. We cleanly link quality attributes to price using observable contract terms established prior to planting. The contract structure allows us to remove price variation driven by potentially endogenous market conditions and isolate variation in price from variation in quality alone.

Second, quality attributes are precisely measured. California's processing-tomato industry has mandatory quality testing by an independent third party, so neither grower nor processor can accidentally or intentionally misstate quality. We observe several individual quality attributes for each field-year observation, which allows us to estimate the impact of weather and climate change on each attribute individually.

Third, our data are unique in that we observe most of the selection and sorting process. Selection in California's processing-tomato industry is relatively small compared to other settings because tomatoes are grown under contract and mechanically harvested and every truckload must be graded at a third-party inspection station. This gives us the opportunity to consider the consequences of selection for the resulting estimates.

We find that extreme weather conditions affect the revenue of growers despite their use of irrigation. An additional 24 hours of exposure to above 40°C temperature decreases yield by 1.8% on average relative to 24 hours at the average temperature (26°C). Quality also declines with exposure to hot temperatures, causing growers to receive a lower price. An additional 24 hours of exposure above 40°C causes prices to fall by 0.2% relative to 24 hours at the average temperature. Taking the effects together, relative to an average temperature, an additional 24 hours of exposure to 40°C temperature decreases revenue by 2%. Failing to account for quality effects would bias downward the effect of exposure to heat on revenue by up to 12%.

We predict the *ceteris paribus* impact of climate change on future processing-tomato production using model parameters that capture adaptive technologies and strategies used today. Accounting for uncertainty in global emissions, climate models, and regression models, we find that climate change will reduce yield and quality without additional adaptation and holding all else equal. Assuming a middle-of-the-road emissions scenario (SSP2–4.5), we predict that by the end of the century, yield will fall 14% due to warming temperatures with a 95% confidence interval of 7% to 31%. We predict losses in quality between 2% and 5% by century's end under a middle-of-the-road emissions scenario. Realized climate change impacts, however, will depend on how industry adapts.

To recover unbiased estimates of temperature and climate change damages, researchers rely on the assumption that observations of quality and yield accurately reflect conditions at harvest. While data are not available to measure the extent, available evidence suggests that selection on quality is likely and perhaps even pervasive in many agricultural settings. Farmworkers hand-harvesting crops—such as berries, lettuce, and other specialty crops—are instructed to pick high-quality products only. For mechanically-harvested crops like corn, soybeans and wheat, low-quality output is routinely withheld from formal markets for use on farm (for example, Kawasaki & Uchida, 2016). This leads to a selected sample and biases observations of both quality and yield—yield will be underestimated and quality overestimated.

To our knowledge, this paper is the first to quantify the consequences of selection on quality in weather and climate change estimates. We replicate the sample-selection problem common in other settings and find that selection leads to statistically significant bias in weather estimates. We find not only that selection biases projections of climate change impacts but that these projections can be misleading. Concerns of selection bias are not limited to papers focusing on quality but apply to *all* studies that estimate yield or revenue effects but do not address selection.

Prior work on the effects of weather and climate change on agricultural production has mostly focused on staple-crop yields (Naylor et al., 2007; Schlenker & Roberts, 2009; Lobell et al., 2011; Tack et al., 2015; Chen et al., 2016; Gammans et al., 2017; Shew et al., 2020; Schmitt et al., 2022). Specialty crops are understudied despite making up 40% of the total value of US crops (USDA NASS, 2017). By focusing on an irrigated specialty crop, we extend a literature that has largely focused on rain-fed staple crops. Irrigated specialty crops have distinct production functions and likely respond differently to weather shocks than rain-fed field crops. Prior work finds that irrigation essentially eliminates the negative effect of extreme heat and climate change on staple crop yields (Shaw et al., 2014; Carter et al., 2016; Tack et al., 2017a; Wing et al., 2021) and agricultural total factor productivity (Liang et al., 2017; Ortiz-Bobea et al., 2018). But we find that in a setting in which irrigation has long been the rule rather than the exception, both yield and quality are affected by exposure to hot temperatures and climate change, leading to lower grower revenue.

2 The Setting

Among all fruits and vegetables, tomatoes rank second in terms of global value of production (FAO, 2023) and second in terms of consumption in the United States (USDA ERS, 2020), behind only potatoes. Tomatoes can be consumed fresh or in processed forms like tomato paste, ketchup, or canned products. They contain essential nutrients such as vitamin E, potassium, and lycopene vital for human health yet frequently underconsumed (Wu et al., 2022).

Tomatoes bound for processing (henceforth "processing tomatoes") are specific varieties, distinct from fresh tomatoes, bred and cultivated to accentuate traits desirable for mechanical harvesting and processing procedures. California's $1 billion processing-tomato industry produces more than 90% of US processing-tomato output (CDFA, 2022). In California, processing tomatoes are typically planted between February and June to enable continuous harvesting from July to October. Tomato plants are primarily cultivated in outdoor fields across the San Joaquin and Sacramento Valleys (CDFA, 2022). They are a warm-season crop and experience average maximum temperatures around 30°C and scant precipitation. Growers irrigate to ensure plants receive enough water and can tolerate high temperatures during the height of summer (Hartz et al., 2008).

Agronomic studies of processing tomatoes (Hartz et al., 2008) find that maximum temperatures between 25°C and 35°C are ideal for vegetative growth, plant development, and fruit set, so long as plants have sufficient water. Hot temperatures without adequate moisture cause tomato plants to become stressed, affecting yield and quality. Temperatures below 10°C slow development and also affect quality. Lobell et al. (2007) find that California's processing-tomato yields benefit from warm temperatures during seedling growth in April but decline when plants are exposed to maximum temperatures above 32°C in June. Marklein et al. (2020) estimate that 34%–87% of land on which tomatoes are grown in California will no longer be suitable by midcentury because summer temperatures will be too hot; they assume hot temperatures translate directly to heat stress, though, so their estimates may be an upper bound. Cammarano et al. (2022) project a decrease in global processing-tomato output by 2050, driven by temperatures rising above the optimal threshold (28°C) in California and Italian growing regions.

We use data from a large tomato processor who buys processing tomatoes under contract from growers in the San Joaquin Valley, Sacramento Valley, and Central Coast region of California. As opposed to purchasing products from a spot market, processors contract with growers to achieve reliable and consistent input quantity and quality. Before the growing season begins, contracts are negotiated between individual processors and the California Tomato Growers Association that represents the interests of all growers. Every year, a base price is negotiated at the start of the growing season that reflects current and expected market conditions. Contracts also establish bonuses and deductions that map observed quality into a price effect. Unlike the base price, we observe very few changes in the schedule of bonuses and deductions across seasons. These contract terms are then offered to growers on a take-it-or-leave-it basis.

As summarized in Table 1, we observe eight measures of quality, of which six are linked to bonuses or deductions. The presence of defects (mold, green tomatoes, worms, material other than tomatoes and limited-use tomatoes) leads to a percentage decrease in a grower's revenue. Additionally, under a brix (soluble solids or sugar content) incentive program, growers receive a bonus (or penalty) if the brix of delivered tomatoes is more (or less) than the county's average. Quality adjustments are a direct function of an individual grower's delivered tomato quality as measured by an independent state inspection station prior to delivery. Finally, the processor prioritizes staggered harvesting and delivery to prevent bottlenecks at processing facilities, and awards producers a bonus for delivering tomatoes either early or late in the season. Quality incentives are economically important—revenue variation from quality is larger than revenue variation from market conditions.

After a truckload undergoes testing at an inspection station, the processing tomatoes onboard are unloaded and washed. Washing removes soil or vines, as well as any juice and pulp from tomatoes that split during transit. Tomatoes are then manually and mechanically sorted in the processing facility. Where possible, the processor uses low-quality tomatoes in their processing. Damaged tomatoes that cannot be used for processing are either composted or used as animal feed.

In most empirical settings, researchers observe only a subset of output that exceeds a minimum quality threshold and is selected to be harvested or graded. Our data are unique in that we observe most of the selection and sorting process. First, there is limited opportunity to selectively harvest and grade processing tomatoes. Unlike many specialty crops, processing

Table 1: Summary of quality measures

Quality attribute	Explanation	Effect on revenue
Brix	Brix is a measure of soluble solids or sugar content	Bonus (penalty) if brix is more (less) than the average for the same variety in the same county
Limited use (LU) percent	Tomatoes that are soft, split, or squashed and have limited processing use	Deduction in proportion to percentage
Material other than tomatoes (MOT) percent	Mainly dirt and sometimes vines	Deduction in proportion to percentage
Green percent	Tomatoes that are unripe	Deduction in proportion to percentage
Mold percent		Deduction in proportion to percentage
Worm percent		Deduction in proportion to percentage
Color score	Red-ripe fruit are given a color score where low scores indicate a better color	No effect
pH	Higher pH usually indicates the fruit is more ripe	No effect

Notes: This table provides some detail on the different measures of quality and how they affect the revenue paid to producers. Additional details were shared with the research team but are held in confidence.

tomatoes are mechanically harvested.[1] Mechanical harvesters sort in the field, and we do not observe what is rejected at harvest. However, according to the processor, there is minimal sorting at harvest. The processor prefers to sort tomatoes at the processing plant because its sorting machines are more accurate than the sorter aboard the harvester. Available empirical evidence supports this claim. For example, we observe truck loads with up to 13% material other than tomatoes, such as dirt or detached stems. We find that 12% of observations have material other than tomatoes above the threshold above which loads could be rejected and turned back to the processor (California Department of Food and Agriculture, 1997). This is consistent with little screening at harvest.

Almost all processing tomatoes grown in California are grown under a contract between a grower and processor (USDA NASS, 2021b). Since operating processing facilities near full capacity is important to profitability, the processor manages the contracts offered to growers to achieve staggered harvesting and delivery. Contracts are written for specific fields, and the processor typically harvests and transports tomatoes to the processing plant. As a consequence, growers have little opportunity to adjust the timing of his/her harvest nor strategically sort prior to delivery.

Further, processing tomatoes in California undergo mandatory grading, including quality measurement. Established in 1987, the California Processing Tomato Inspection Program requires an independent state inspection station to grade every truckload of tomatoes prior to delivery. These quality measurements are used to determine each grower's revenue.

In sum, the institutional setting means quality and yield observations accurately reflect conditions in the field; our analysis is consequently less susceptible to selection problems common in other settings. The setting also allows us to quantify the magnitude of biases introduced if we simulate selection.

With proprietary data, there is a trade-off between internal and external validity. Using data from a large processor offers a level of detail not available in public data. These data come not from surveys nor field trials but from administrative records of every field contracting with the processor between 2011 and 2021. Our data are at the field-year level and contain a range of information about fields such that we can observe and control for field-specific factors. Greater detail enhances internal validity but may come at the cost of

[1] See Just & Chern (1980) for details on the introduction and widespread adoption of mechanical harvesters in California's processing-tomato industry during the 1960s.

limited external validity.

One concern is that the hundreds of growers in this proprietary dataset are not representative of the broader processing-tomato industry. Several factors mitigate this concern. The fields in our sample are geographically dispersed across 18 counties and closely match the spatial distribution of output in California (Figure 1). Field-level yields averaged at the county level are 10% higher than county yields reported by NASS, and the series are highly correlated. Nevertheless, it is always possible that there is selection on unobservables that affects the type of grower we observe and thus the external validity of our results.

Another concern is that our results are relevant only for processing tomatoes. However, as noted above, virtually all agricultural producers are paid a price that depends on quality and face similar incentives to processing-tomato growers, albeit with different contract provisions. The practice of selective harvesting based on a quality threshold is pervasive across crops. As a result, findings from this case study offer important insight into the effects of weather on yield and quality in other agricultural settings.

3 Data

3.1 Field-Level Dataset

We use data on all processing tomatoes grown and sold under contract to a large processor in California. Observations are at the field-year level between 2011 and 2021 ($n = 13,292$) and include information on field acreage, tomato variety, total tons, yield, quality attributes, and the latitude and longitude of the field centroid. Quality and tonnage observations come from mandatory testing administered by the Processing Tomato Advisory Board. The data also include information about growing practices, including planting and harvesting dates, irrigation technology, and the crop previously planted on a field. Field-level observations are linked to an unbalanced panel of 480 growers and 272 grower groups. The grower-group identifier links growers within the same network or organization.[2] We also observe the annual pricing terms negotiated between the processor and the California Tomato Growers Association at the start of the growing season.

[2] An example of a grower group is four siblings dividing a family farm. Each sibling would have a distinct grower ID, and the four would share a common grower-group ID.

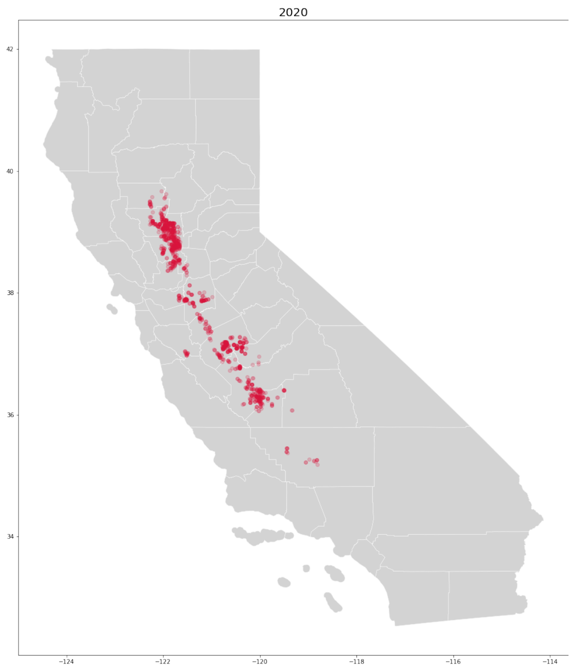

(a) Location of fields contracted with processor in 2020

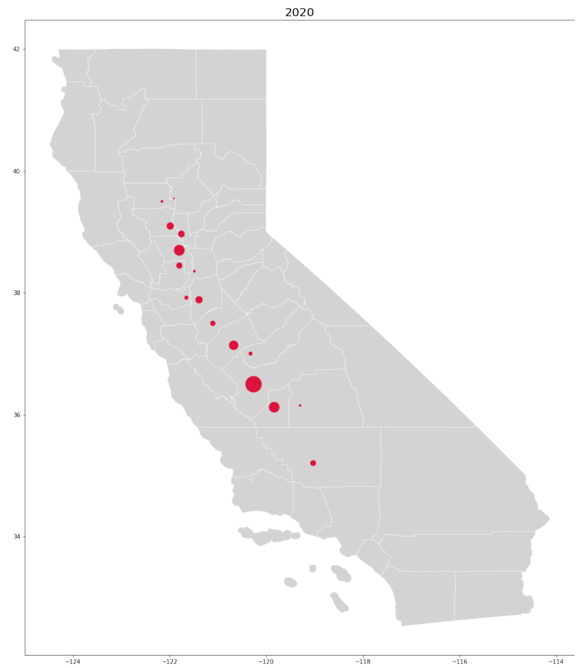

(b) California output in 2020, by county (USDA NASS, 2021a)

Figure 1: California processing-tomato output

Notes: These maps show the spatial distribution of processing-tomato output in California according to two data sources: proprietary data from a tomato processor, and public data from USDA NASS (2021a).

3.2 Outcomes

Yield for field i in year t is defined as total tons per acre: $\text{yield}_{it} = \frac{\text{total tons}_{it}}{\text{acres}_{it}}$. *Total tons* is all tomatoes harvested from a field irrespective of quality. We exclude material other than tomatoes, such as dirt or vines, from the total tons used to calculate yield.

Next, we calculate price for each field-year observation using observed quality and the schedule of quality bonuses and deductions established at the beginning of the growing season. We isolate variation in price driven by quality by applying the quality bonuses and deductions to the 11-year-average base price. This removes common price movements driven by market shocks while preserving common and individual quality shocks. This *quality-adjusted price* is not the price received by growers but, effectively, a quality index with weights equal to each quality attribute's effect on price. Our method of calculating price avoids potential simultaneity bias in our estimates of price and revenue effects (see Section 4 for more details). The quality-adjusted price is calculated as follows:

$$\text{price}_{it}^{\text{quality adjust}} = \left(\overline{\text{base price}} + \text{bonus}_{it}\right) \times (1 - \text{deducts}_{it}) \tag{1}$$

Here, $\overline{\text{base price}}$ is the average base price in the 11-year sample and $\textit{deducts}_{it}$ and $\textit{bonus}_{it}$ are adjustments that depend on observed quality and date of delivery of tomatoes from field i in year t. Deductions and bonuses affect price differently. The presence of defects lead to a proportional deduction from the price, whereas bonuses (in dollars per ton) are added to the price.

Finally, we estimate field-level revenue per acre (henceforth "revenue") by multiplying quality-adjusted price by total tons and dividing by acreage as follows:

$$\text{revenue}_{it} = \frac{\text{total tons}_{it}}{\text{acres}_{it}} \times \text{price}_{it}^{\text{quality adjust}} \tag{2}$$

3.3 Weather Data

We obtain weather data from PRISM (2020), which publishes daily temperature and precipitation data interpolated to 4 km grids for the whole time span. We match weather data to each field-level observation by mapping the field centroids to PRISM grid cells. In Section 4.1, we explain how we translate daily observations of temperatures into measures of temperature exposure for each field-year observation.

3.4 Control Variables

We also gather data on several controls. We source information on tomato varieties from AgSeeds (2020), which includes key attributes and use categories for each of the 159 varieties in the processor dataset. Finally, we match each field to its major soil type in the National Cooperative Soil Survey (NRCS USDA, 2020).

Table 2: Summary statistics

	units	mean	sd	min	max
Area	acres	57.21	43.18	0.3	323.2
Growing days	no.	133.96	9.43	96.0	175.0
Yield	tons/acre	52.61	13.32	6.7	99.0
Quality attributes					
Brix		5.08	0.49	3.5	7.2
LU percent		1.42	1.14	0.0	13.1
MOT percent		1.74	1.28	0.0	15.1
Green percent		3.13	2.40	0.0	24.3
Mold percent		1.68	1.91	0.0	27.5
Worm percent		0.00	0.01	0.0	0.5
Color score		21.04	1.69	13.3	34.4
pH		4.41	0.09	2.9	4.8
Weather					
Average minimum temperature	° C	13.77	1.07	9.8	17.6
Average maximum temperature	° C	31.31	1.63	24.5	35.9
Total precipitation	mm	22.45	25.73	0.0	198.9
Soil type					
Alluvium	prop.	0.96	0.20	0.0	1.0
Eolian	prop.	0.00	0.04	0.0	1.0
Organic material	prop.	0.03	0.17	0.0	1.0
Lacustrine	prop.	0.00	0.03	0.0	1.0
Residuum	prop.	0.01	0.08	0.0	1.0
Irrigation technology					
Drip irrigation	prop.	0.76	0.43	0.0	1.0
Furrow irrigation	prop.	0.12	0.33	0.0	1.0
Missing irrigation tech.	prop.	0.10	0.30	0.0	1.0
Sprinkler irrigation	prop.	0.02	0.14	0.0	1.0
Varietal attributes					
Extended field storage variety	prop.	0.53	0.50	0.0	1.0
Tomato spotted wilt resistant	prop.	0.47	0.50	0.0	1.0
Fusarium Wilt resistant	prop.	0.21	0.40	0.0	1.0
Powdery Mildew resistant	prop.	0.04	0.21	0.0	1.0
Fusarium Crown Rot resistant	prop.	0.01	0.08	0.0	1.0
Bacterial Spot resistant	prop.	0.00	0.05	0.0	1.0
High solids	prop.	0.05	0.22	0.0	1.0
High yield	prop.	0.05	0.22	0.0	1.0
Early maturing	prop.	0.13	0.34	0.0	1.0
Thin consistency	prop.	0.13	0.34	0.0	1.0
Intermediate consistency	prop.	0.25	0.43	0.0	1.0
Thick consistency	prop.	0.59	0.49	0.0	1.0
Pear-shaped	prop.	0.01	0.08	0.0	1.0

Notes: Summary statistics on all field observations from 2011 to 2021 ($n = 13,292$).

4 Methods

The aim is to estimate the effect of weather on processing-tomato yield, quality, and revenue. We focus on the effect of temperature exposure and include precipitation as a control since precipitation during the growing season is scant and growers control the amount of water applied through irrigation.

Since we are interested in the direct and indirect effects of temperature exposure, we are careful not to introduce bad controls—variables that themselves are outcome variables. A key example in our setting is irrigation volume, which is a function of temperature and also affects the outcome variable. Were irrigation volume included, the coefficient on temperature exposure could be biased because some of the explanatory power of temperature might be incorrectly attributed to irrigation volume.

We take an off-the-shelf econometric approach to emphasize that results are driven by differences in focus and setting rather than differences in methodology. We follow the approach proposed by Schlenker & Roberts (2009) and adopted by Gammans et al. (2017) and Shew et al. (2020), among others. Ortiz-Bobea (2021) provides a comprehensive summary. We estimate the following equation:

$$y_{it} = \int_h f(h)\phi_{it}(h)d(h) + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{3}$$

Here, y_{it} is a log-transformed outcome variable (yield, quality, or revenue) in field i in year t, $\alpha_{g(i)}$ is a grower fixed effect, and $\psi(t)$ is a quadratic year trend. The first term characterizes the relationship between temperature exposure and the outcome variable, where $f(h)$ is the marginal effect of temperature h and $\phi_{it}(h)$ is the growing-season density of exposure at h for field i in year t. This continuous representation is not tractable for estimation but can be approximated using the restricted cubic-spline specification detailed in Section 4.2. Field-year-specific control variables z_{it} include variety-specific attributes (extended field storage, various disease-resistance traits, high amount of solids, high yield, early maturing, thin consistency, intermediate consistency, thick consistency, and pear shaped), irrigation technology (drip, sprinkler, furrow), soil type (alluvium, eolian, organic material, lacustrine, and residuum), growing-season precipitation, a dummy for planting week, and the difference between actual growing days and estimated growing days specified by the seed manufacturer.

As with any annual crop, growers can implicitly influence their expected weather by varying the planting date. Specifically, tomatoes planted earlier in the growing season are expected to be exposed to cooler temperatures than tomatoes planted later. This implies that weather is endogenous and coefficients on temperature exposure may be biased. We include dummies for planting week of year to account for endogeneity of weather. Failing to control for planting date biases estimates of effects on quality, but the results for yield and revenue are largely unchanged (see Appendix A).

The error term ϵ_{it} is likely heteroskedastic, spatially correlated, and temporally correlated among similar growers over time. We use heteroskedasticity-robust standard errors clustered by grower group[3] and county-by-year. We cluster at the grower-group level to account for possible temporal dependence among growers within the same grower group. Clustering by year would allow for spatial correlation across all observations in the same year, but we observe only 11 years of data which is too few clusters to provide accurate statistical inference (Cameron et al., 2011). Instead, we cluster at the county-by-year level to account for spatial correlation within a county, a relatively large geographic area. Results are effectively unchanged if we cluster at the irrigation-district-by-year level instead of county-by-year level. Results are also robust to using spatial heteroskedasticity- and autocorrelation-consistent errors that allow for spatial correlation between nearby fields and serial correlation in panel data (see Appendix B).

We include grower fixed effects $\alpha_{g(i)}$ and individual growers can be associated with multiple fields per year; on average, each grower is associated with 27 field-year observations. This controls for time-invariant, grower-specific factors that may be related to outcome or explanatory variables. Our preferred specification uses grower fixed effects and controls for observable field characteristics, as regular crop rotation results in an unbalanced panel of field-year observations. Unobservable, time-invariant field characteristics are not a concern—results are robust to using field fixed effects in place of grower fixed effects (see Appendix C). Results are also robust to replacing quadratic year trends with (a) a linear year trend and (b) county-specific quadratic year trends.

We do not use year fixed effects, as they would absorb much of the useful variation in temperature exposure used to identify the effects of interest. Inclusion of state-by-year fixed effects (equivalent to year fixed effects in our setting, as we only observe a single state) in Deschênes & Greenstone (2007) is critiqued by Fisher et al. (2012) because "state-by-year absorb almost all variation and the identification rests on very slim margins, so even

[3]Recall that the grower-group identifier links growers within the same network or organization.

small amounts of measurement error will be greatly amplified." When excluding year fixed effects, we observe considerable variation in temperature exposure (Appendix D shows the distribution of temperature exposure across counties).

Since we do not include year fixed effects, a reasonable concern is that our estimates of price and revenue effects suffer from simultaneity bias. Simultaneity may be caused by equilibrium in the tomato market: when quantity is low, price is high. We avoid this issue by applying the quality adjustments to an average price, which removes variation in price (and revenue) that stems from market conditions and preserves variation caused by quality, as detailed in Section 3. In other words, we do not need year fixed effects to control for price or price expectations—their effect has already been removed. Results using contemporaneous prices are shown in Appendix E.

4.1 Estimating Temperature Exposure

We translate daily observations of minimum and maximum temperatures into a measure of growing-season temperature exposure for each field-year observation. Each field's growing season starts on the day the field was planted and ends on the final day of the field's harvest. For each day of each field's growing season, we estimate how many hours are spent in 1°C temperature intervals by fitting a sinusoidal curve between each day's minimum and maximum temperature. We then sum over days to estimate the number of days spent in each 1°C temperature interval during the entire growing season. The result is x_{it}, a 1-by-J vector of temperature exposure for field i during the growing season in year t, where J is the number of temperature bins. In our setting, we bin temperatures from 5°C to 41°C, so $J = 37$.[4] The vector is thus

$$x_{it} = \begin{pmatrix} x_{it,5} & x_{it,6} & \dots & x_{it,40} \end{pmatrix}, \tag{4}$$

where $x_{it,j}$ is the number of days spent between j°C and $(j+1)$°C during the growing season in year t.

This approach has several advantages. First, it addresses the empirical challenge of mixed frequency between regressor and outcome variables. We have many daily observations of

[4] Temperatures range from -1°C to 45°C. We aggregate temperature exposure below 5°C and above 41°C to avoid bins with little exposure.

minimum and maximum temperatures to match with one annual observation of an outcome variable. Averaging daily temperatures across the growing season would mask differences in exposure to extreme temperatures. The second advantage of this approach is that it preserves the temperature distribution. This allows us to uncover the marginal effect of exposure to different temperatures.

4.2 Restricted Cubic-Spline Specification

Next, we choose a functional form to characterize the relationship between outcome variables and temperature exposure. Midrange temperatures are thought to be ideal for yield and quality of processing tomatoes, but yield and quality may be reduced by hot or cool temperatures if exposure occurs during key stages of a plant's growth cycle (Hartz et al., 2008). Thus the relationship between temperature exposure and outcome variables is nonlinear.

To capture nonlinearity in the outcome variable's response to temperature, we estimate a restricted cubic-spline model (otherwise known as a natural cubic spline). The restricted cubic-spline model has become popular because it offers several benefits over alternative methods for estimating nonlinear temperature effects (Berry et al., 2014; D'Agostino & Schlenker, 2016; Ortiz-Bobea et al., 2019; Blanc & Schlenker, 2020; Bucheli et al., 2022). First, it offers smooth, parsimonious semiparametric estimation without needing to define critical temperature thresholds. Second, it imposes the restriction that its tails (that is, before the first knot and after the last knot) are linear. This reduces overfitting in the data-sparse tails of the temperature distribution, an issue with the polynomial and cubic-spline functional forms. We estimate a piecewise-linear degree-day model as a robustness check. Overall, results from the two specifications—a piecewise-linear degree-day model and a restricted cubic-spline model—are consistent in magnitude and significance (see Appendix F).

We identify $K = 4$ temperatures that split the distribution of temperature exposure by interval into quintiles. This accounts for the fact that less time is spent at extreme temperatures. Unlike in other models, knot placement does not strongly influence the cubic-spline results because the marginal effect of exposure is allowed to vary smoothly between knots. In Appendix G, we vary the number and placement of knots and find that results are largely unchanged.

We then introduce a basis matrix B and a vector of coefficients Γ. B is J-by-P, while Γ

is P-by-1, where P is the number of temperature parameters to be estimated and is directly related to the number of knots K. A restricted cubic spline with $K = 4$ knots results in $P = 3$, which is smaller than the number of temperature-exposure bins $J = 37$. This is an advantage of the spline model: it reduces the dimensionality while still allowing for flexible semiparametric estimation. The derivation of the basis matrix B that corresponds to the restricted cubic spline is shown in Appendix H.

Under these assumptions, we can write Equation 3 as follows:

$$y_{it} = x_{it} B\Gamma + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{5}$$

Stacking n observations across fields and years gives us Equation 5 in matrix notation:

$$Y = XB\Gamma + \delta Z + \alpha + \psi + \epsilon \tag{6}$$

Here, X is an n-by-J matrix of temperature exposure, and Y, δZ, α, ψ, and ϵ are n-by-1 vectors of outcomes, controls, grower fixed effects, quadratic time trends, and errors, respectively.

After estimation, we recover the marginal effect of temperature exposure evaluated at each interval. We multiply the vector of estimated coefficients $\hat{\Gamma}$ by the corresponding B matrix. The resulting J-by-1 vector $\hat{\beta}$ is the marginal effect of one additional day in each temperature bin $j = 1, \ldots, J$:

$$\underset{J\times 1}{\hat{\beta}} = \underset{J\times P}{B} \times \underset{P\times 1}{\hat{\Gamma}} \tag{7}$$

Last, we derive an estimate of the variance-covariance matrix for $\hat{\beta}$:

$$\underset{J\times J}{\widehat{var}(\hat{\beta})} = \underset{J\times P}{B} \times \underset{P\times P}{\widehat{var}(\hat{\Gamma})} \times \underset{P\times J}{B'} \tag{8}$$

5 Results

Figure 2 reports results for the effects of temperature exposure on our three key outcome variables: yield, quality, and revenue. In each figure, the top graph shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to

24 hours at 26°C, which represents the average temperature. The 95% confidence intervals account for the possibility of heteroskedasticity, spatial correlation, and temporal correlation in the errors. The omitted category, 26°C, has no confidence interval. The gray vertical lines show the positions of the knots. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

Regarding yield, we find that the optimal temperature is around 28°C. Exposure to temperatures above 35°C leads to significantly lower yields. An additional 24 hours of exposure to above 40°C temperature decreases yield by 1.8% on average relative to 24 hours at the average temperature. Exposure to temperatures below 10°C causes a small but significant decline in yield—0.8%—relative to 24 hours at the average temperature.

We isolate variation in price driven by quality by applying quality adjustments to the 11-year-average base price. Regarding quality, we find an optimal temperature around 20°C. Quality declines with exposure to hot conditions. An additional 24 hours of exposure above 40°C causes quality to drop by around 0.2% relative to 24 hours at the average temperature. Results for specific quality defects and bonuses can be found in Appendix I, Figures A.17 and A.18. Defects—specifically, limited-use tomatoes, material other than tomatoes[5], green tomatoes, and mold—all increase with exposure to hot temperatures, although a lack of precision in the estimates means we cannot rule out null effects. Exposure to cool temperatures relative to average temperatures leads to fewer limited-use tomatoes and higher quality. The effect of temperature on the solids bonus is imprecise, but the point estimate declines with exposure to hot temperatures.

Revenue is maximized with exposure to temperatures around 27°C. An additional 24 hours of exposure to 40°C temperature decreases revenue by 2% compared to 24 hours at the average temperature. This is expected since both yield and quality decline under hot conditions. Exposure to temperatures below 10°C causes a smaller but still significant decrease in revenue—almost 1%—relative to 24 hours at the average temperature.

Appendix K shows the estimated effects of the control variables. Precipitation reduces yield (and therefore revenue), but the magnitude is relatively small—a one-standard-deviation increase in precipitation decreases yield by 0.2% on average. Fields with furrow or

[5] A priori, we did not expect temperature to influence the presence of material other than tomatoes (MOT). As a robustness check, we estimate the effect of temperature on quality excluding MOT in Appendix J. The magnitude of the quality effect excluding MOT decreases slightly but the overall result is unchanged.

sprinkler irrigation techniques are associated with lower yields on average relative to fields using drip irrigation (the dropped category). Soil type and varietal characteristics are also associated with yield and quality. For example, an early variety, one that requires fewer days to reach maturity and therefore has a shorter season, is associated with lower yields but higher quality on average. The quadratic function of the year trend is insignificant in all our regression models, suggesting that yield and quality were largely stationary over the course of our 11-year sample. This confirms visual evidence in Appendix D that shows no obvious trends in yield or quality.

5.1 Decomposition

We find that temperature exposure significantly affects both yield and quality. However, the relative importance of each pathway is unclear a priori. Here, we decompose the effect of temperature exposure on revenue per acre (total effect) into the effect on revenue driven by yield (yield effect) and the effect driven by quality (quality effect) as shown in Equation 9. This allows us to answer two questions. What is the relative importance of the yield and quality effects? And, perhaps more importantly, would the estimates of the revenue effect be biased if quality were omitted?

$$\text{total effect} = \text{yield effect} + \text{quality effect} \tag{9}$$

Since we use log-transformed variables, these effects are equal to the effect of exposure estimated above. As shown in Figure 3, while the yield effect dominates grower revenue, quality also plays an important role. Without access to data on output quality, a researcher can only recover the yield effect: an additional 24 hours of exposure to 40°C temperature decreases revenue by 1.8% compared to 24 hours at the average temperature. This underestimates the effect of exposure on revenue by up to 0.2 percentage points, or 12% of the point estimate at 40°C. Failing to account for quality's effect on revenue biases estimates of temperature on revenue.

5.2 Selection

In almost all published work, to recover unbiased estimates of temperature and climate change damages, researchers implicitly assume observations of quality and yield accurately

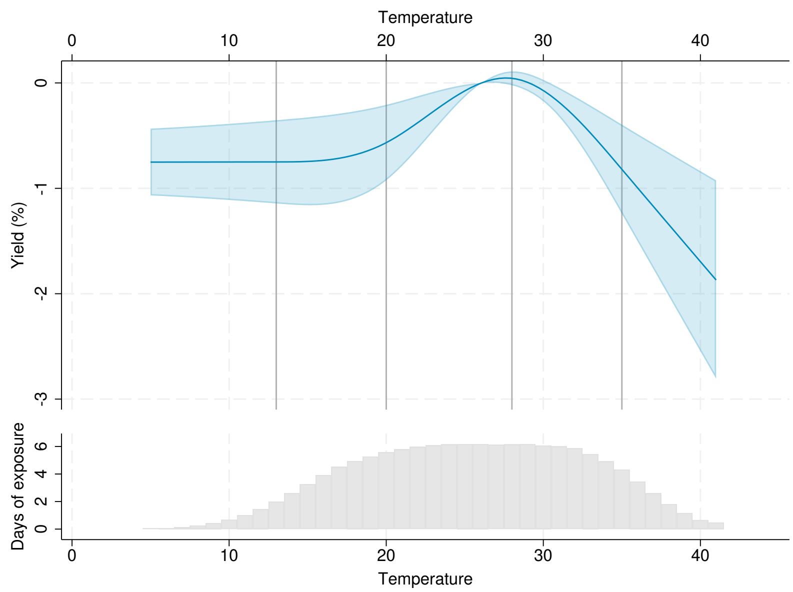

(a) Yield

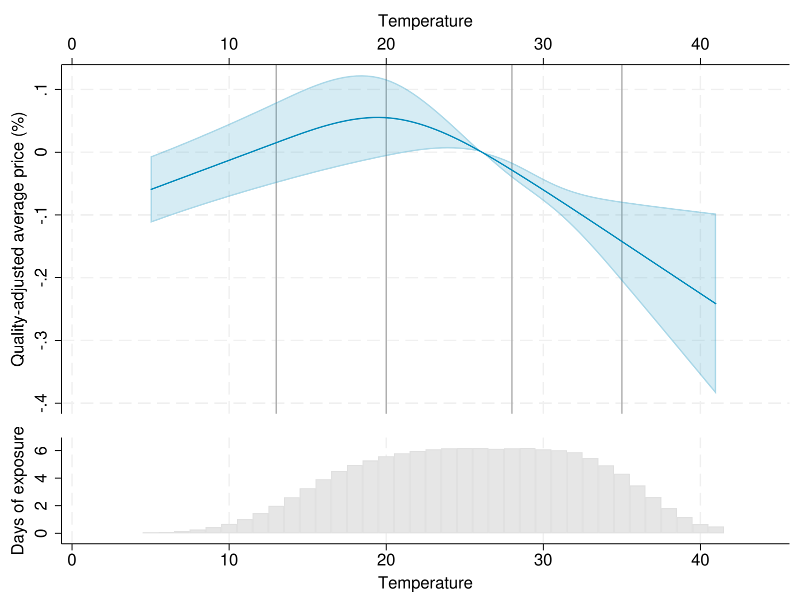

(b) Quality

Figure 2: Restricted cubic-spline results

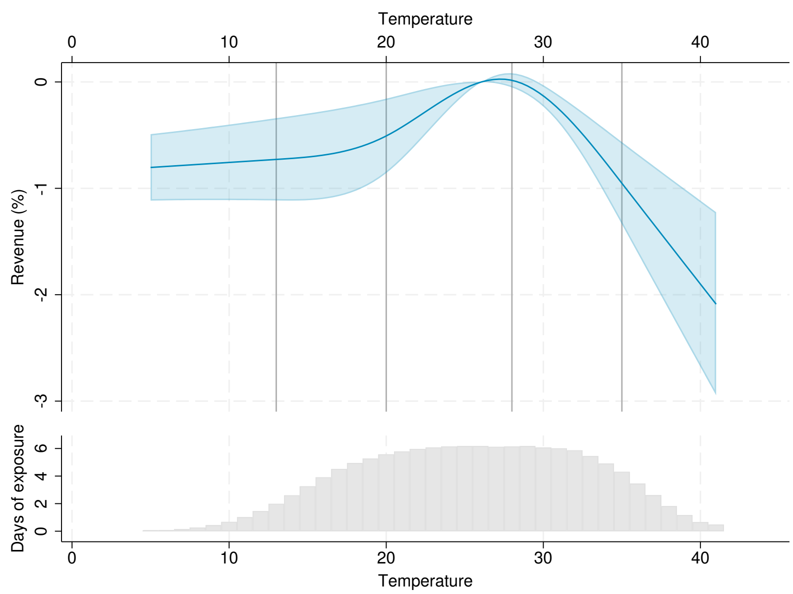

(c) Revenue per acre

Figure 2: Restricted cubic-spline results

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

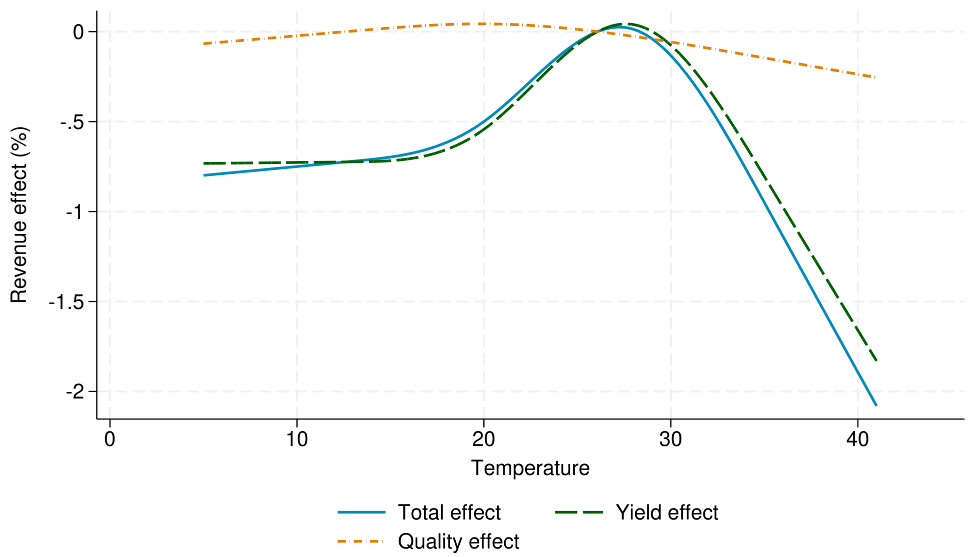

Figure 3: Decomposing the effect of exposure on revenue into the yield effect and quality effect

Notes: Each line shows the effect of an additional 24 hours in a given temperature interval on the outcome relative to 24 hours at 26°C. The total-effect line shows the effect of temperature exposure on revenue through both the yield and quality pathways. The yield-effect and quality-effect lines show the contributions of yield and quality to the total revenue effect.

reflect harvest conditions. For many agricultural products, data are available only for the subset of output that producers select to be graded. This leads to a selected sample and biases observations of both quality and yield—for example, if only high-quality products are graded, yield will be underestimated and quality overestimated.

A benefit of our setting is that selection is minimal. Recall from Section 2 that observations of quality and yield are close to those in the field because of mandatory grading as well as contracting and harvesting practices. This gives us the opportunity to consider the consequences of selection on quality as experienced in other settings for the resulting estimates.

Under the California processing tomato inspection program marketing order, tomatoes may be rejected by a processor if defects exceed a specific limit (California Department of Food and Agriculture, 1997). We use these inspection program limits to simulate selection in our setting. We assume defective tomatoes in excess of the limits are withheld from sale and are therefore unobserved. Yield under selection is artificially reduced because it is calculated using a smaller tonnage than what is actually harvested from the field. Quality under selection is artificially improved because some tonnage with defects is no longer observed. Figure 4 compares the actual distribution of defects to its distribution with selection.

A priori the sign of any bias is unclear; however, our first hypothesis is that the effect of exposure to high temperatures on quality will be biased toward zero under selection. This follows from the result that high temperatures reduce quality and from the assumption that lower-quality products are being withheld. Our second prediction is that the negative effect on quality will be incorrectly assigned to yield, causing the negative yield effect to increase in magnitude.

We compare estimates from our preferred specification with estimates using observations with added selection. Consistent with our hypothesis, the estimated effect of exposure on yield is biased upward by up to 10%. The estimated effect of high temperatures on quality is biased toward zero under selection by up to 85%. Selection causes statistically significant differences in the estimated relationships between temperature exposure and both yield and quality (see Appendix L). For revenue, however, the upward bias in yield offsets the downward bias in quality so that the effect of exposure to high temperatures on revenue is unchanged. Our results suggest that estimates will be biased in settings with selection but the magnitude will depend on (a) the actual effect of weather on quality and (b) how much

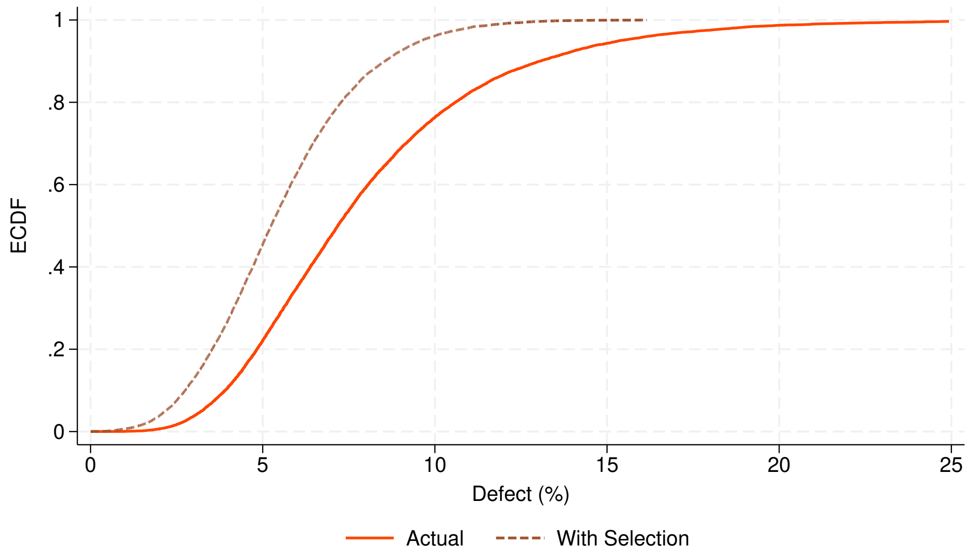

Figure 4: Empirical cumulative distribution function of defects

Notes: The actual CDF shows the distribution of defects observed across fields and years in our sample, in which selection is known to be minimal. The CDF with selection shows the distribution of defects when we simulate selection on quality as experienced in other settings.

selection is occurring.

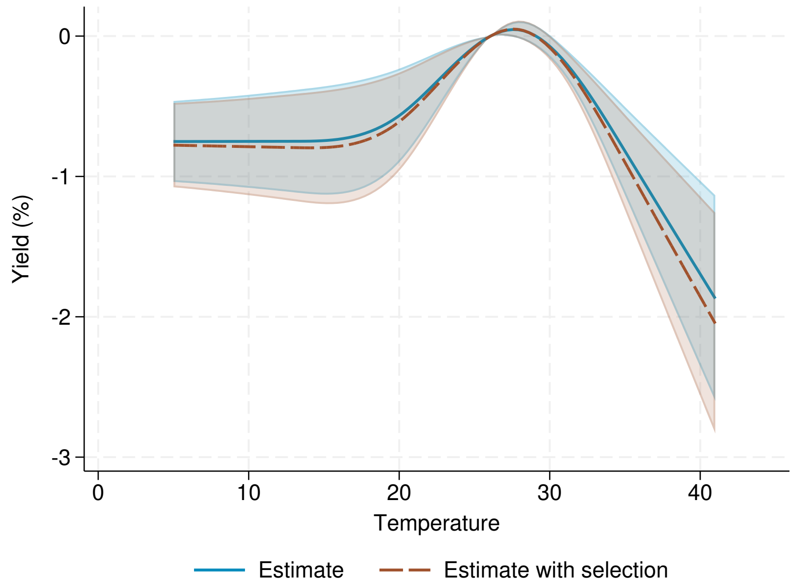

(a) Yield

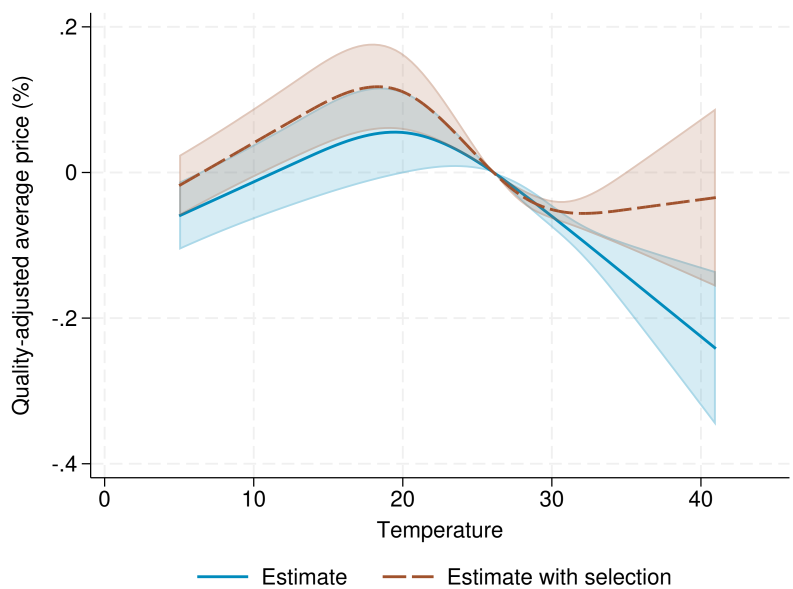

(b) Quality

Figure 5: Selection bias

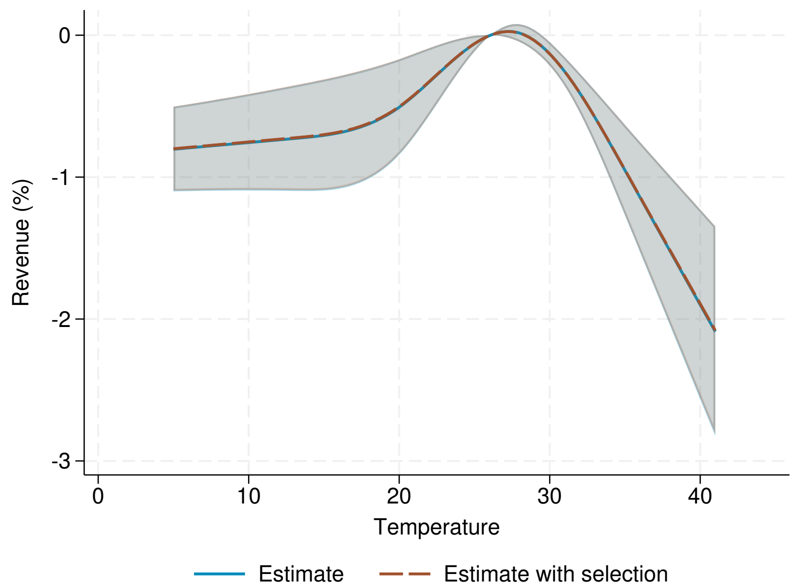

(c) Revenue per acre

Figure 5: Selection bias

Notes: These graphs show the effect of an additional 24 hours in a given temperature interval on the outcome variables relative to 24 hours at 26°C. The estimate line shows results using our preferred specification. The estimate with selection shows results using yield and quality observations with added selection as experienced in other settings. The difference between the two lines shows the degree of bias caused by selection.

6 Projected Impacts of Climate Change

The results above suggest processing tomatoes are susceptible to extreme temperatures. We now ask: how will climate change affect the production of processing tomatoes? We predict yield, quality, and revenue outcomes as though tomato plants are exposed to projected temperatures in climate models. We then compare predicted outcomes using realized weather with predicted outcomes using future weather to estimate the *ceteris paribus* impacts of climate change. While rain-fed staple crops have been extensively studied, this is among the first efforts to forecast the effect of climate change on the yield and quality of specialty crops despite their representing 40% of the total value of US crops (USDA NASS, 2017).

In estimating the impact of climate change on economic outcomes, uncertainty stems from several sources. First, there is statistical uncertainty in the historical relationship between weather variables and the outcomes of interest. Second, future global emissions of greenhouse gases are uncertain. Third, conditional on a particular pathway of global emissions, there is uncertainty about how that level of emissions will change the climate of a particular location. While climate models generally agree on the overall direction of temperature and precipitation trends, there is disagreement on the exact magnitudes—that is, there is climate-model uncertainty. Our estimates of climate change impacts account for these three sources of uncertainty.

As with any projection, there are limits on how much uncertainty we can incorporate into our impact estimates. It is difficult to know how producers, processors, seed manufacturers, and other industry participants will respond to future changes in climate. For example, the processing-tomato industry may become more resilient to hotter temperatures through developing and adopting varieties better suited to a warmer climate. On the other hand, decreasing availability of water for irrigation (Hayhoe et al., 2004; Elliott et al., 2014) may make processing-tomato plants more vulnerable to warming temperatures than historical relationships imply. Modeling the potential effects of these factors (and others[6]) on future processing-tomato production is beyond the scope of this paper. We instead assume the modeled relationship between temperature exposure and processing-tomato output that captures adaptive technologies and strategies used during our sample period will continue.

Following best practices outlined in (Burke et al., 2015), our approach to projecting

[6] We also do not account for potential yield benefits from the predicted increase in atmospheric carbon dioxide concentrations (Rangaswamy et al., 2021; Cheng et al., 2022).

climate change impacts proceeds as follows. First, we select four global climate models included in Coupled Model Intercomparison Project Phase 6 (CMIP6) and used by the Intergovernmental Panel on Climate Change in its latest assessment report (IPCC, 2023): Access-CM2, HadGEM3-GC31-LL, EC-Earth3, and EC-Earth3-Veg. These models best capture relevant aspects of California's climate (Krantz et al., 2021) and were statistically downscaled by Pierce et al. (2023) to a 3 km resolution and daily time step (available on Cal-Adapt, 2023). To account for uncertain future global emissions, we present results for two emissions scenarios in CMIP6: SSP2–4.5, a middle-of-the-road global emissions scenario, and SSP5–8.5, a very high global emissions scenario.

To account for climate-model uncertainty, we use four climate models and collect multiple climate projections, or ensembles, with varying baseline conditions from each model when available. This yields a total of 10 projections for each emissions scenario. Each projection carries equal weight in determining the estimated effect—known as the model-democracy method commonly used by climate scientists (Burke et al., 2015).

To address statistical uncertainty in our modelling of the historical relationship between weather and outcomes, we use a wild-cluster bootstrap procedure. This method preserves dependence across clusters in resampled data (Cameron et al., 2011).

To implement the wild-cluster bootstrap, we first obtain residuals $\hat{\epsilon}$ and predicted coefficients after estimating our preferred regression model in Equation 6. Next, we generate 1,000 bootstrap samples denoted by $*$ and indexed by b for each grower-group-and-year cluster g:

$$Y_g^{*b} = X_g B\hat{\Gamma} + \hat{\delta} Z + \hat{\alpha} + \hat{\psi} + \epsilon_g^{*b}, \qquad \epsilon_g^{*b} = w_g^{*b}\hat{\epsilon}_g \tag{10}$$

Here, ϵ_g^{*b} are the randomly reshuffled residuals $\hat{\epsilon}_g$ from the same grower-group-and-year cluster multiplied by a wild weight w_g^{*b} drawn from the Rademacher distribution (-1 or 1 with equal probability). We then regress Y^{*b} on X to obtain $\hat{\Gamma}^{*b}$. Finally, we predict $\hat{Y}_{itp}^{*b}$ by replacing x_{it}, temperature exposure experienced in field i during year t with x_{ip}, temperature exposure projected to be experienced in field i in year p at midcentury ($p = 2041, \ldots 2050$) or end of century ($p = 2091, \ldots, 2100$) under a particular combination of climate model projection and emissions scenario. We keep all controls and fixed effects at their original year-t levels. This estimates the outcomes as though the field was exposed to temperatures from future year p instead of the actual temperatures experienced in year t.

We estimate the *ceteris paribus* climate change impact for each bootstrap replication as the difference between the predicted outcome using actual temperatures and predicted outcomes using projected temperatures from midcentury (2041–50) and end of century (2091–2100). For each emissions scenario, we stack the 1,000 bootstrap estimates from each of the 10 projections into a vector of 10,000 impact estimates. We calculate a 95% confidence interval accounting for both statistical and climate uncertainty by taking the 2.5th and 97.5th percentiles of the vector of impact estimates.

Perhaps unsurprisingly given the results above, we find that climate change leads to economic damages for processing-tomato growers absent additional adaptation and all else equal, as shown in Figure 6. By midcentury, yield is predicted to be significantly lower than its 2011–21 levels. We estimate median yield losses of 6% or 9% depending on the emissions scenario. The median effects on quality by midcentury are negative in both emissions scenarios, but under the middle-of-the-road emissions scenario the 95% confidence interval includes zero. By the end of the century, yield and quality are expected to both decline significantly by 14% and 3% respectively under a middle-of-the-road emissions scenario, with even larger losses predicted under a very high emissions scenario (53% and 9% respectively). The projected impacts of climate change on revenue are provided in Appendix M.

When estimating these climate change impacts, we made two relatively strong assumptions: no additional adaptation and holding all else equal. Realized climate change impacts will depend on adaptation or maladaptation in the processing-tomato industry. While fitted time trends imply yield and quality were stationary between 2011 and 2021, longer time horizons may see secular increases in processing-tomato production via technology improvements. The locations where processing tomatoes are grown may change. We also do not account for the potential yield boost from rising atmospheric carbon dioxide concentrations. Modeling how these factors may influence future processing-tomato production is beyond the scope of our paper.

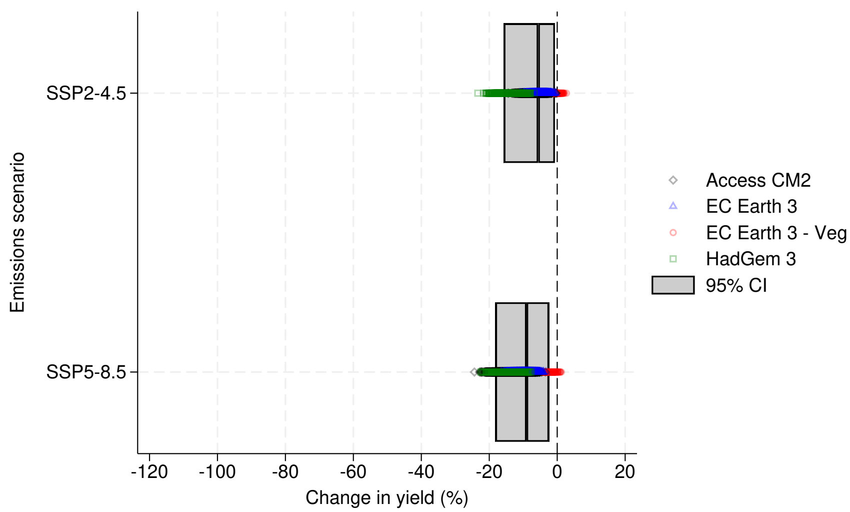

(a) Yield, midcentury

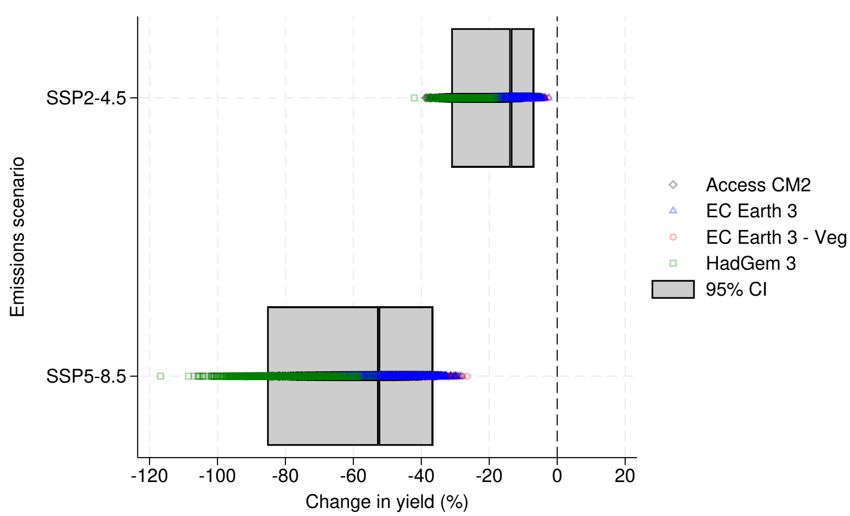

(b) Yield, end of century

Figure 6: Projection of climate impacts by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

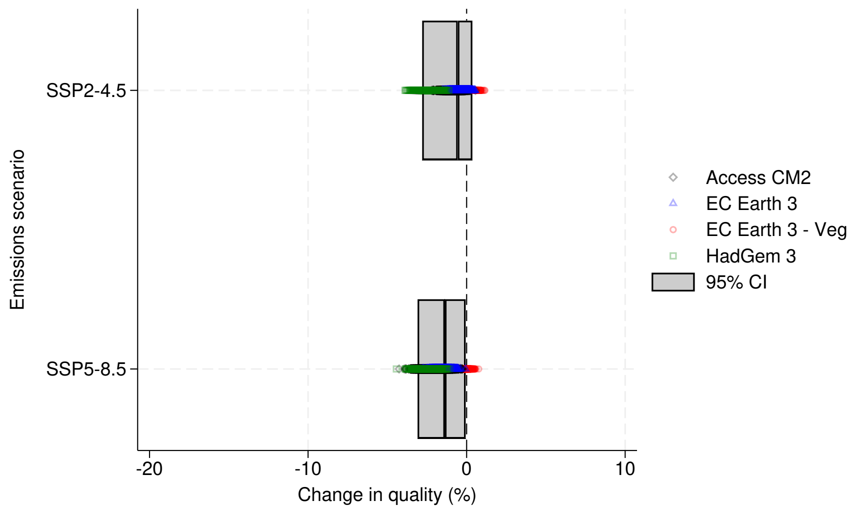

(c) Quality, midcentury

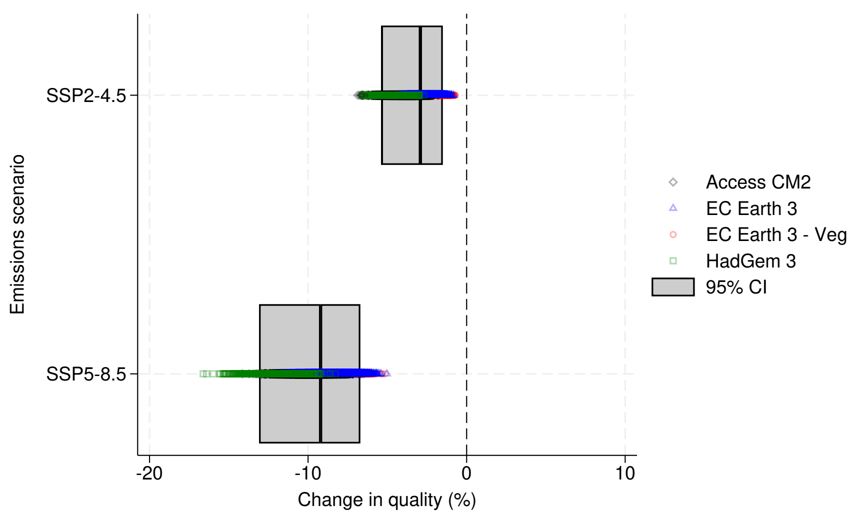

(d) Quality, end of century

Figure 6: Projection of climate impacts by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

Notes: These graphs show the estimated impact of climate change on the outcome variables by midcentury and end of century, assuming no additional adaptation and all else equal. Each point is an estimate of the projected impact derived from a single combination of projection, emissions scenario, and wild-cluster bootstrap replication. The thick black lines represent the median impact estimate, and the shaded gray areas represent the 95% confidence intervals that account for statistical and climate uncertainty.

We also estimate the effect of a uniform increase in daily temperatures by 1°C, 2°C, and 3°C, shown in Table 3. This simplified model of rising temperatures circumvents uncertainty from future emissions scenarios and climate projection models. However, we still implement the wild-cluster bootstrap procedure described above to capture statistical uncertainty in the relationship between temperature and the outcomes in the 95% confidence intervals. We find that a uniform increase in growing-season temperature causes a statistically significant and meaningful decline in yield, quality, and revenue.

Table 3: Estimated percentage change in yield, quality and revenue caused by a uniform increase in daily temperatures

	Yield	Quality	Revenue
+1°C	-2.94	-1.24	-4.22
	(-2.25, -3.69)	(-1.09, -1.39)	(-3.49, -4.92)
+2°C	-7.22	-2.63	-10.05
	(-5.74, -8.87)	(-2.32, -2.95)	(-8.43, -11.54)
+3°C	-12.95	-4.15	-17.62
	(-10.52, -15.64)	(-3.65, -4.66)	(-14.97, -20.14)

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header. 95% confidence interval (in parentheses) are estimated using a wild-cluster bootstrap procedure with 1,000 simulations.

In Section 5.2, we replicated the sample-selection problem, common in other settings, in which quality losses are incorrectly assigned to yield losses. We found estimates of the damages from exposure to hot temperatures on quality are artificially reduced by selection bias, while yield damages from exposure to hot temperatures are artificially increased. Here, we ask: does selection also bias estimates of climate change impacts?

We repeat the climate change–impact estimation procedure above using yield and quality observations with added selection as experienced in other settings. Figure 7 compares the median climate change impacts with and without selection (the full results with selection are available in Appendix N).

We find that selection biases the magnitude of climate change estimates and, in some cases, the bias is large enough that estimates are in the wrong direction. Climate change damages to yields are increased by selection: the median estimate increases from -14%

without selection to -17% with selection by the end of the century under a middle-of-the-road global emissions scenario. Selection also biases quality results to the extent that damages are incorrectly identified as benefits: the median estimate without selection is -1.5% compared to 1% with selection by midcentury under a high global emissions scenario. Climate change–impact estimates could be misleading when researchers rely on data that suffer from selection on quality.

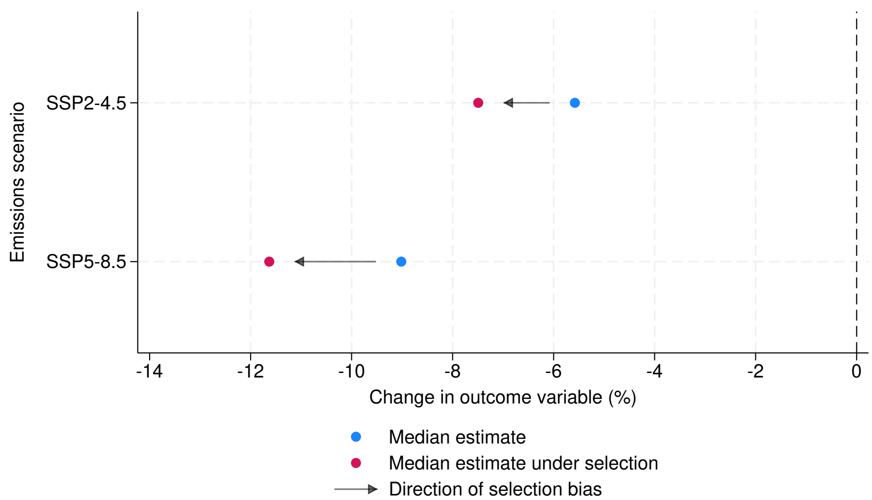

(a) Yield, midcentury

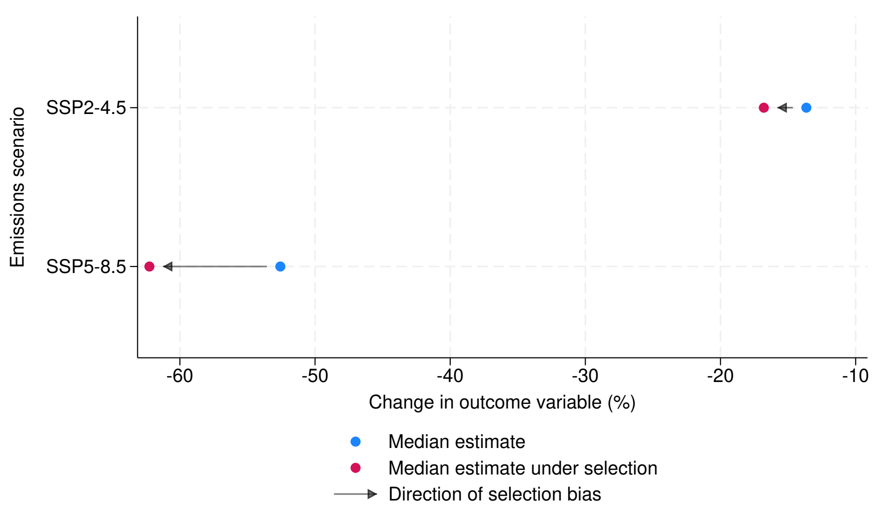

(b) Yield, end of century

Figure 7: Selection bias in median climate impacts by midcentury (2041–50) and end of century (2091–2100)

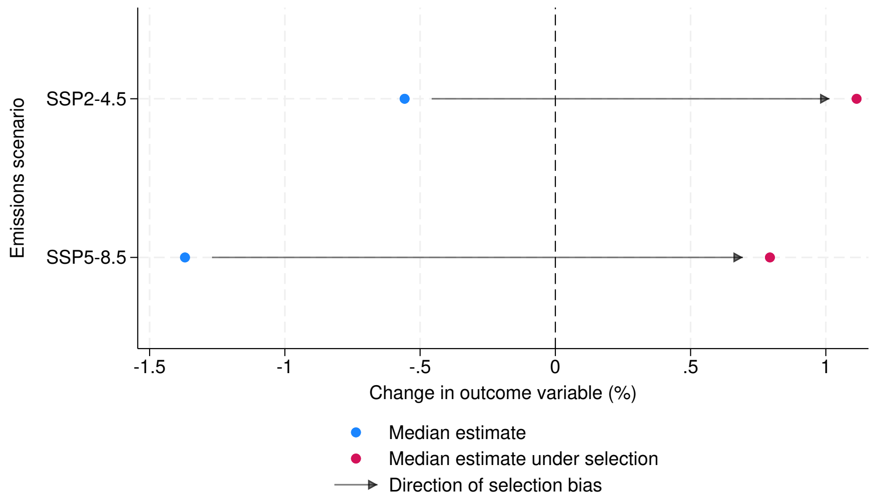

(c) Quality, midcentury

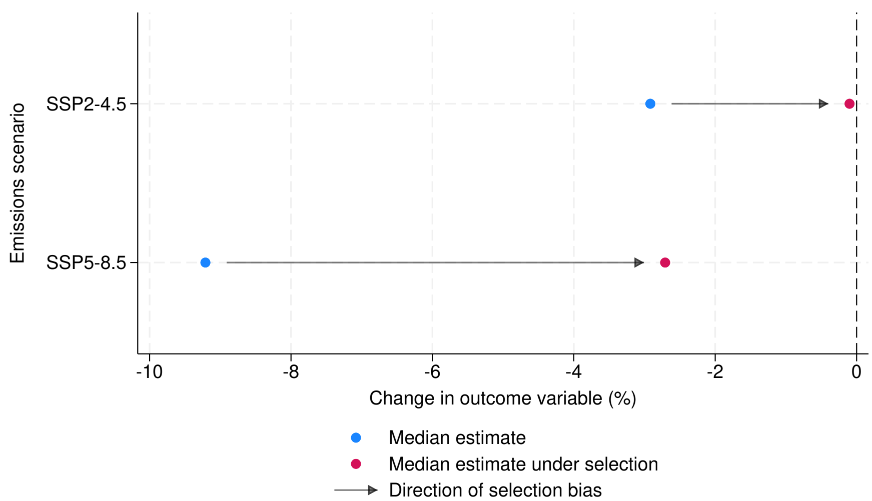

(d) Quality, end of century

Figure 7: Selection bias in median climate impacts by midcentury (2041–50) and end of century (2091–2100)

Notes: These graphs show the median effect of climate change on the outcome variables by midcentury and end of century, assuming no additional adaptation and all else equal. The median-estimate point shows results using our preferred specification. The median estimate under selection shows results using yield and quality observations with added selection as experienced in other settings. The arrow connecting the two points shows the direction of bias caused by selection.

7 Discussion and Conclusion

Studies of the impacts of climate change on agriculture generally focus on yield (see Ortiz-Bobea, 2021) or farmland values (for example, Mendelsohn et al., 1994; Fezzi & Bateman, 2015; Hendricks, 2018; Bareille & Chakir, 2023). More recent work explores alternative effects of weather shocks and climate change on agriculture such as pest and disease pressure (Kawasaki, 2023) and planting and harvesting decisions (Cui, 2020; Cui & Xie, 2022), but with a few notable exceptions, ignore quality. Quality matters for essentially every agricultural product because of its role in contractual arrangements and price determination. We used novel data from a large tomato processor to study the effects of temperature exposure on yield, quality, and grower revenue. We found that exposure to hot temperatures reduces grower revenue through two channels: yield and quality. Failing to account for quality's effect biases estimates of temperatures on revenue by up to 12%. Our uniquely detailed data gives us a complete and precise picture of how weather and climate change affect product quality.

Prior work uses observational data to analyze effects of weather and climate on the quality of different agricultural products, including rice and wheat (Kawasaki & Lichtenberg, 2014; Kawasaki & Uchida, 2016; Kawasaki, 2019), apples (Dalhaus et al., 2020), peanuts (Ramsey et al., 2020), tobacco (Ramsey & Rejesus, 2021, but the consequences of climate change were not considered), and wine (see Ashenfelter & Storchmann, 2016). Of these papers, Kawasaki & Uchida (2016), Kawasaki (2019), Dalhaus et al. (2020), and Ramsey et al. (2020) link changes in quality to grower revenue to quantify the economic consequences of weather and climate change. Dalhaus et al. (2020) infer quality from unexplained differences in price. Kawasaki & Uchida (2016) and Kawasaki (2019) use quality grades, which inhibits their ability to analyze individual quality attributes. Ramsey et al. (2020) observe one aspect of peanut quality, kernel size, and proxy for its effect on price using value formulas from the Commodity Credit Corporation's loan rates. It is unclear whether this captures actual market pricing and all relevant aspects of quality. In our setting, quality attributes are precisely measured by a third party, rather than inferred, and directly linked to price using a schedule of bonuses and deductions established prior to planting.

In addition, studying the effects of climate change on agricultural production with observational data presents identification challenges. Typically, data are available only for the subset of output that growers choose to market because it exceeds an implicit or explicit quality threshold.[7] For example, in Japan, around two-thirds of rice output undergoes the costly process of being graded while the remaining lower-quality output is withheld for self-consumption or sale in informal markets (such as sales to local households or for animal feed) (Kawasaki & Uchida, 2016). In the United States, observations of corn sold on formal markets would miss between 6%–15% of output used on farm (USDA ERS & NASS, 2023). Observational data for hand-harvested crops, such as berries, stone fruit, apples, and leafy greens, likely also suffer from selection. Harvesting guidelines (for example, instructions to only pick high-quality produce), adequate access to labor, and how labor is paid can all affect the observed quality of hand-harvested crops (for example, Hill & Beatty, 2024).

The practice of screening out low-quality products introduces selection bias into weather and climate change estimates. Quality effects are biased when low-quality products are not observed. Yield estimates will also be biased if quantities are measured after farmers screen out low-quality products. In this situation, quality effects will be falsely attributed to yield, and estimates of the effect of weather on both yield and quality will be biased. For example, selection might explain why prior work finds that wet conditions toward the very end of the season appear detrimental to crop yields (Ortiz-Bobea et al., 2019). Wet conditions around harvest typically result in quality problems in grain that may cause growers to withhold low-quality grain from the market. If this withheld grain is not counted toward output totals, an outside observer will incorrectly ascribe the effects of wet conditions around harvest to yield instead of quality. Concerns about selection are not limited to papers focusing on quality but are present in *all* studies estimating the effects of weather on yield or revenue that do not address selection.

Processing tomatoes in California are almost always grown under contract, mechanically harvested, and graded at an independent state inspection station. These institutional factors mean that our analysis is unlikely to suffer from selection bias. It also gave us the opportunity to estimate selection akin to that experienced in other settings and quantify the magnitude of the bias. We found that screening out low-quality products biases the negative effect of hot temperatures on quality toward zero and that this negative effect on quality is incorrectly assigned to yield. We also found that error caused by selection is magnified when projecting

[7] An exception is data from field trials such as those used by Ramsey et al. (2020), Tack et al. (2015) and Tack et al. (2017b), which are less likely to suffer from selection bias but may be less representative of commercial growing practices.

the impact of future warming under climate change. We found that reductions in quality due to climate change are underestimated in the presence of selection and that, in some cases, climate change is incorrectly predicted to improve quality. Reductions in yield from warming under climate change are biased upward by selection. These results illustrate the consequences of data limitations that should be carefully considered in all empirical work.

Finally, continued warming predicted under climate change is a cause for concern. Based on retrospective results, we see that growers cannot fully mitigate damages caused by extreme heat, indicating that they will be susceptible to harm from the continued warming predicted as the climate changes. Absent additional adaptation, we predict climate change will cause both yield and quality of California's processing tomatoes to decline by the end of the century. This contrasts with previous findings that irrigation can mitigate the effects of heat on staple crop yields in settings in which irrigation is the exception rather than the norm (Shaw et al., 2014; Carter et al., 2016; Tack et al., 2017a; Wing et al., 2021). Our setting differs in that production relies on irrigation. Our results suggest that irrigated agriculture is susceptible to climate change and that irrigation's mediating effect on heat may be short-lived. The results above emphasize the need for investment in research into and development of heat-tolerant varieties and related technologies.

A limitation of our research is its focus on a single processor in one agricultural industry. While the specificity of our setting potentially limits the external validity of our results, our study benefits from detailed and reliable observations of field-level outcomes that capture farmers' optimizing behavior. This provides new insights—often hidden in analyses of county-level averages—into the effect of weather and climate on individual commercial agricultural producers.

A Endogeneity in Weather

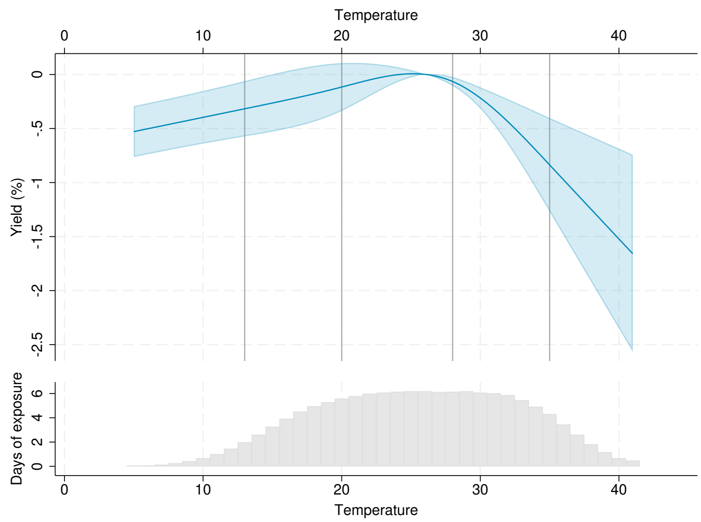

(a) Yield

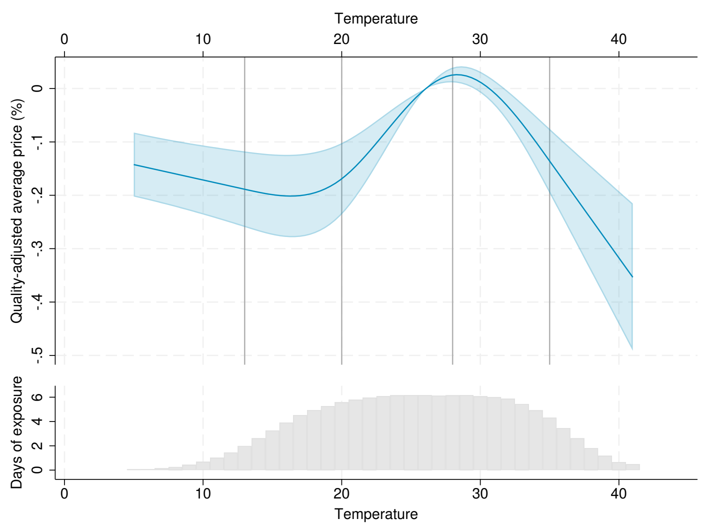

(b) Quality

Figure A.1: Results without a planting date control

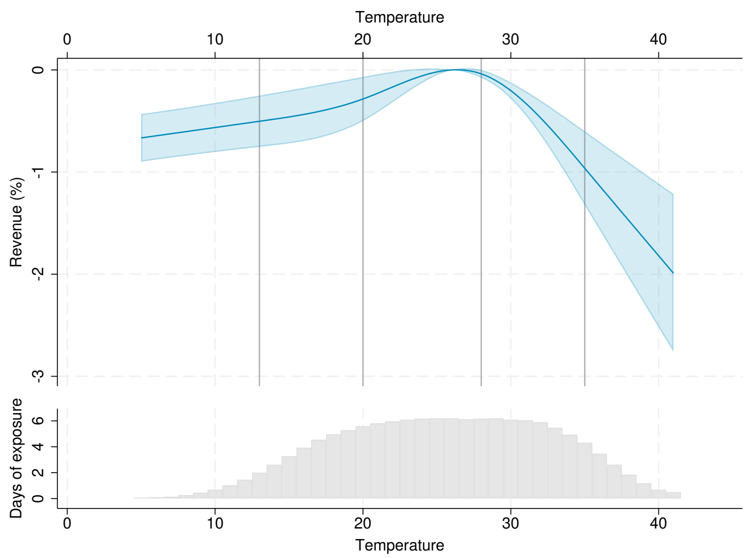

(c) Revenue per acre

Figure A.1: Results without a planting date control

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

B Alternative Methods for Estimating the Standard Errors

B.1 Clustering by Irrigation Districts

In our preferred specification, we cluster standard errors by grower group to account for temporal dependence and county-by-year to account for spatial dependence. Here, we show the consequences of replacing county-by-year clusters with irrigation-district-by-year clusters.

Processing tomato growers rely on water for irrigation from two sources: groundwater and surface water, of which most comes from surface water. Irrigation districts are organizations that source surface water, mostly from state and federal water projects, and distribute it to the farmers within their boundaries. Using irrigation district boundaries from the California Department of Water Resources via California State Geoportal (2023), we match tomato fields to the irrigation district they fall within. Irrigation districts are often smaller than counties—the fields in our sample are located within 19 counties and 174 irrigation districts.

Here, we show results that use standard errors clustered by grower group and irrigation district by year. The standard errors are slightly smaller and the 95% confidence intervals slightly narrower in Figure A.2 than our main results using county-by-year clusters. This is expected when clustering on a smaller spatial unit. In our preferred specification we allow for dependence between all fields in the same county, whereas here we are enforcing independence between fields within the same county but different irrigation districts.

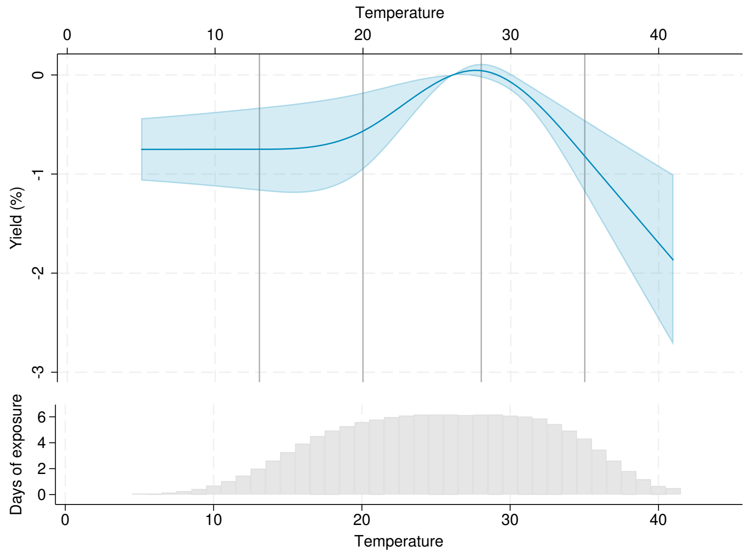

(a) Yield

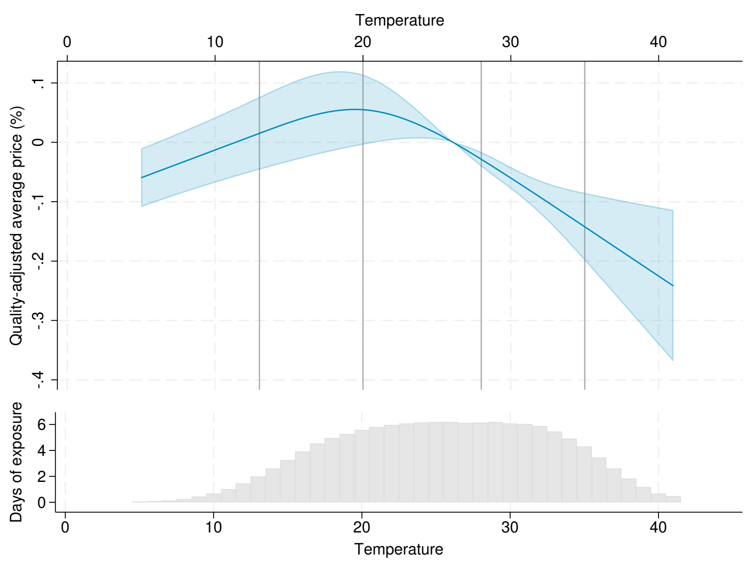

(b) Quality

Figure A.2: Results using standard errors clustered by grower group and irrigation district

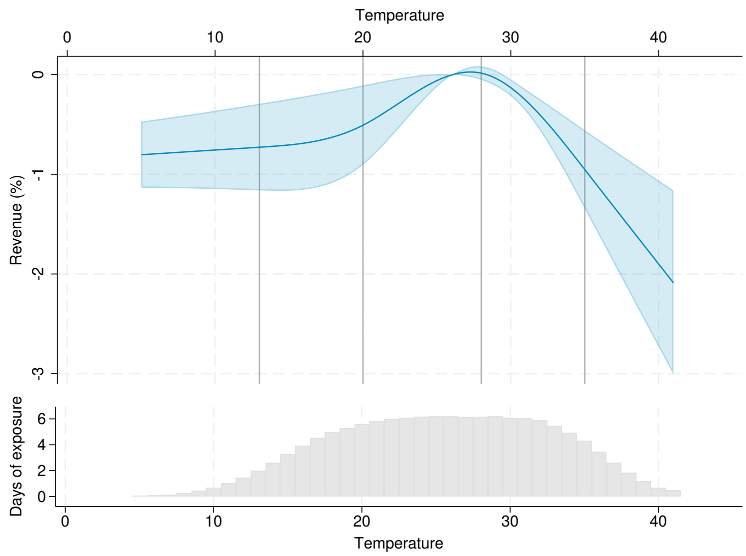

(c) Revenue per acre

Figure A.2: Results using standard errors clustered by grower group and irrigation district

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

B.2 Spatial Heteroscedasticity and Autocorrelation Consistent Errors

Our preferred specification clusters standard errors by grower group and county by year to account for the possibility of heteroskedasticity, spatial correlation, and temporal correlation in the errors. An alternative method to correct for possible dependence in standard errors is to estimate spatial heteroscedasticity and autocorrelation consistent (HAC) errors that allow for spatial correlation and serial correlation in panel data (Conley, 1999). Using code from Hsiang (2010), we allow for spatial correlation for field observations that are within 200km (124 miles) of each other. The correlation between observations is assumed to decay linearly with distance.

Results are robust to using spatial HAC errors, however the appropriateness of using this method on an unbalanced panel remains an open question.

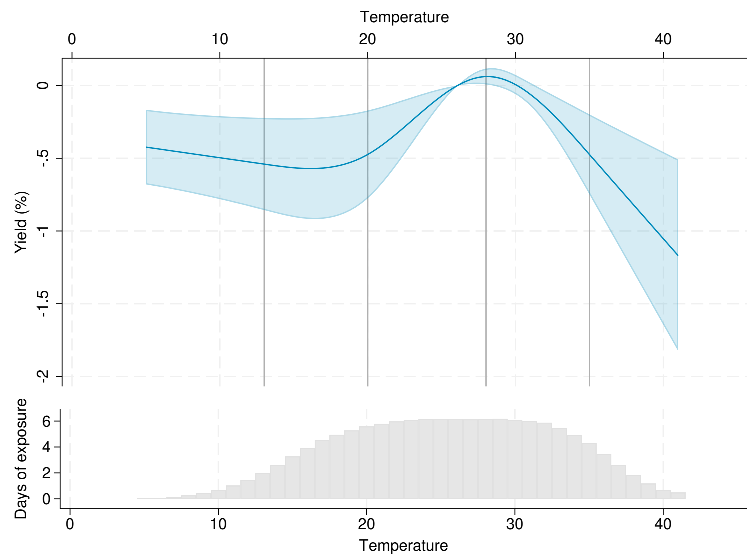

(a) Yield

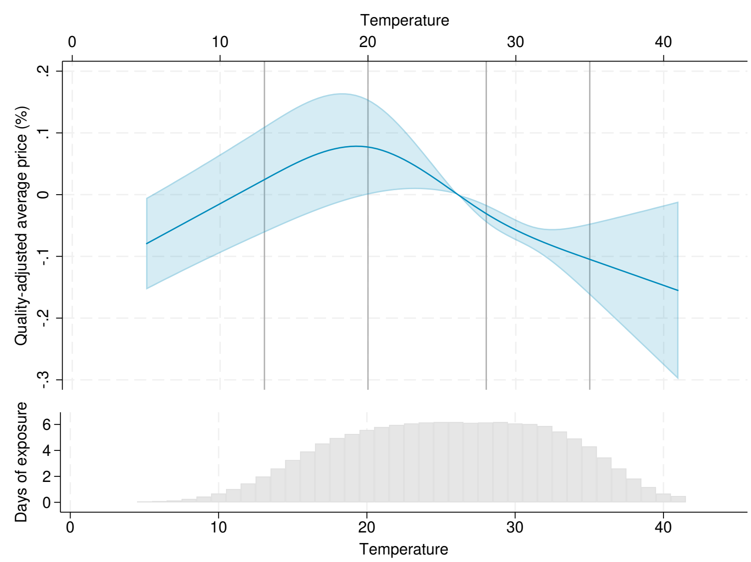

(b) Quality

Figure A.3: Results using spatial HAC errors

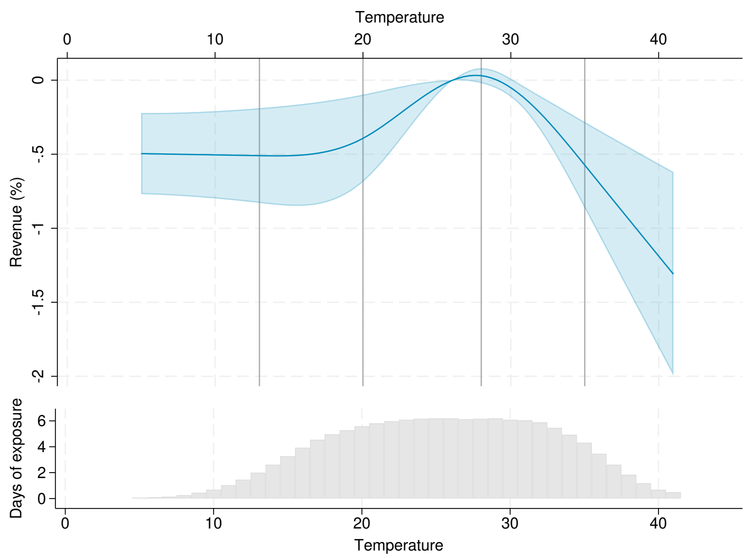

(c) Revenue per acre

Figure A.3: Results using spatial HAC errors

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

B.3 Wild Cluster Bootstrap Errors

In this robustness check, we use a wild cluster bootstrap to calculate the 95% confidence interval around the estimated effect of temperature exposure on yield, quality, and revenue following Cameron et al. (2008). The procedure is as follows:

1. Estimate the main model in Equation 6 for log-transformed yield, quality, and revenue: $Y = XB\Gamma + \delta Z + \alpha + \psi + \epsilon$

   (a) Save coefficient estimates

   (b) Save residuals from yield and quality regressions

2. Generate 1,000 bootstrap samples denoted by * and indexed by b. For each bootstrap simulation, generate wild-bootstrap residual ϵ_g^{*b}:

   (a) For yield and quality, $\epsilon_g^{*b} = w_g^{*b}\hat{\epsilon}_g$ where $\hat{\epsilon}_g$ are the randomly reshuffled residuals from the same grower-group-and-year cluster multiplied by a wild weight w_g^{*b} drawn from the Rademacher distribution (-1 or 1 with equal probability)

   (b) For revenue, $\epsilon_{g,revenue}^{*b} = w_g^{*b}\left(\epsilon_{g,yield}^{*b} + \epsilon_{g,quality}^{*b}\right)$, which accounts for joint uncertainty in yield and quality on revenue

3. Using coefficient estimates from step 1 and wild-bootstrap residuals from step 2, generate the wild-bootstrap outcome variables $Y_g^{*b} = X_g B\hat{\Gamma} + \hat{\delta}Z + \hat{\alpha} + \hat{\psi} + \epsilon_g^{*b}$

4. Regress Y_g^{*b} on X to obtain $\hat{\Gamma}^{*b}$

5. Recover the marginal effect of temperature exposure at each interval: $\hat{\beta}^{*b} = B \times \hat{\Gamma}^{*b}$

6. Recover the variance of the marginal effect of temperature exposure: $s_{\hat{\beta}^{*b}} = \sqrt{\widehat{var}\left(\hat{\beta}^{*b}\right)} = \sqrt{B \times \widehat{var}\left(\hat{\Gamma}^{*b}\right) \times B}$

7. Create a list of candidate null hypotheses $H_0 : \beta = \beta_0^n$ where n indexes a null hypothesis

8. Generate Wald statistic for each bootstrap and null hypothesis: $W_n^{*b} = \left(\hat{\beta}^{*b} - \beta_0^n\right) / s_{\hat{\beta}^{*b}}$

9. For each null hypothesis, calculate its p-value as the proportion of times that $|W_n| > |W_n^{*b}|$ for $b = 1, \ldots, B$

10. Calculate the 95% confidence interval is the set of null hypotheses β_0^n for which $p \geq 0.05$

Overall, the 95% confidence intervals using wild cluster bootstrap are similar in magnitude to those in our preferred results. Unlike the smooth and monotonic intervals produced using other methods, the intervals produced using the wild cluster bootstrap are jagged. This is an artifact of the bootstrap method using combinations of discrete null hypotheses and temperature intervals.

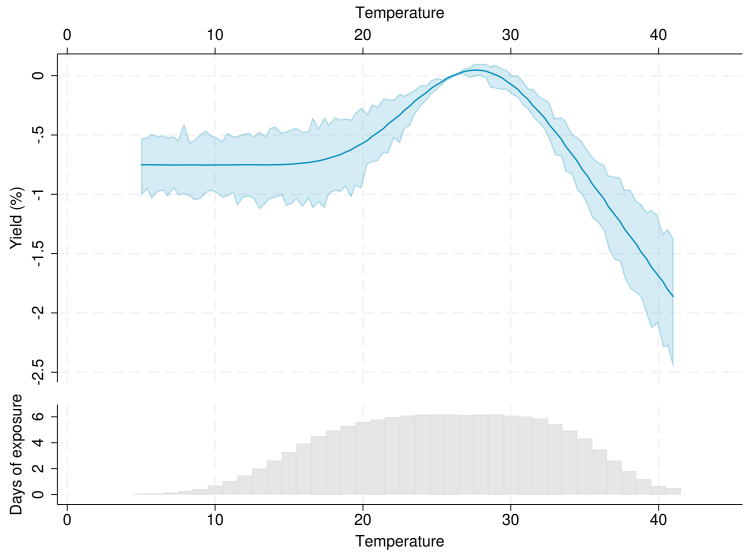

(a) Yield

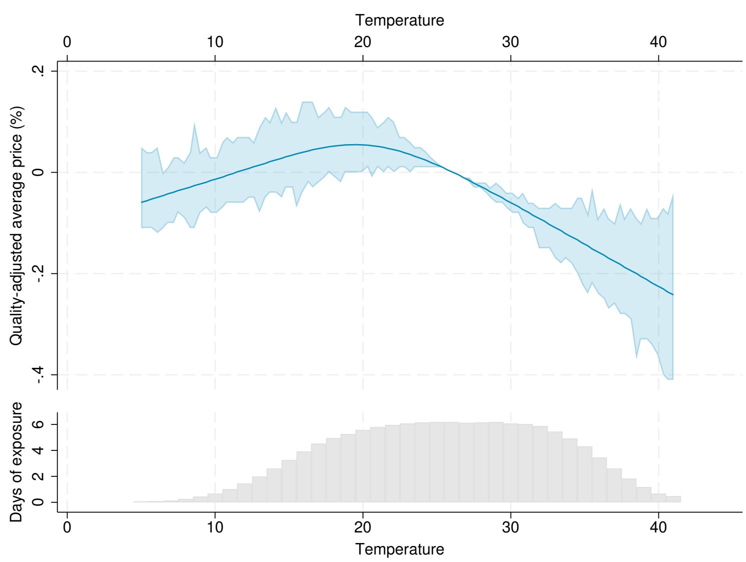

(b) Quality

Figure A.4: Results using wild cluster bootstrap errors

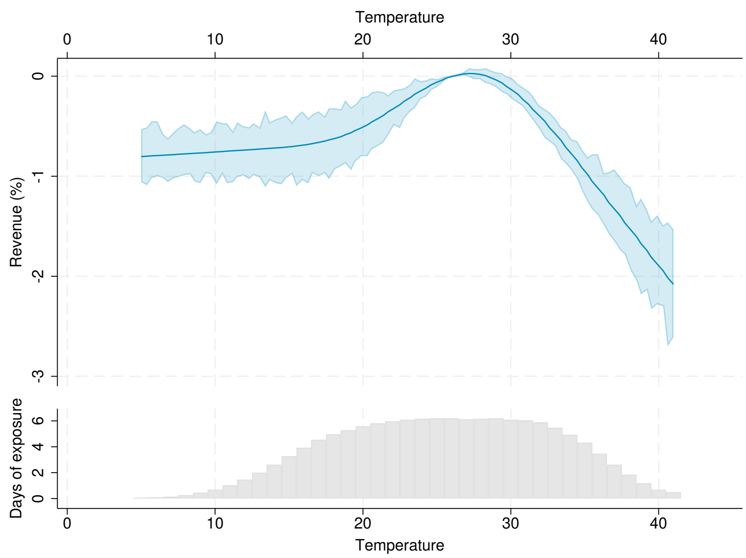

(c) Revenue per acre

Figure A.4: Results using wild cluster bootstrap errors

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

C Robustness Checks

Our preferred specification uses grower fixed effects. Grower fixed effects capture time-invariant characteristics of the growers and characteristics of their respective fields that are both time-invariant and common across fields. One might be concerned that fields differ in ways that are correlated with weather, which would introduce omitted variables bias into our estimation. To alleviate this concern, we estimate Equation 14 using field fixed effects instead of grower fixed effects. Some fields do not appear multiple times in our sample because of crop rotation. We drop around 30% of field-year observations because they do not make a field-level panel. The results from this estimation are similar to those from the estimation using grower fixed effects.

We also replace the quadratic year trend with a (a) linear year trend, and (b) county-specific quadratic year trends. The results are robust to the choice of functional form for the time trend.

We include the number of days between planting and harvesting to control for the season length. The effect of exposure to hot temperatures is virtually unchanged. Exposure to cool temperatures is estimated to improve yield, quality and revenue although the effects are statistically indistinguishable from zero.

Our preferred specification includes control variables like variety attributes and irrigation technology that are decided before the growing season begins and are therefore unlikely to be influenced by growing-season temperature exposure. However, it is possible that these control variables are bad controls if seasonal forecasts influence these decisions by growers and/or the processor. As a robustness check, we remove field-year controls (except precipitation and planting week of year) and find that the results are largely unchanged.

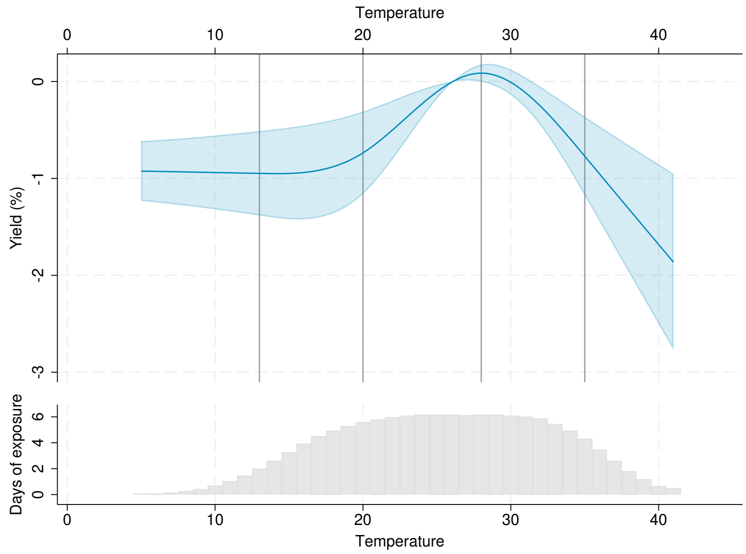

(a) Restricted spline with field fixed effects

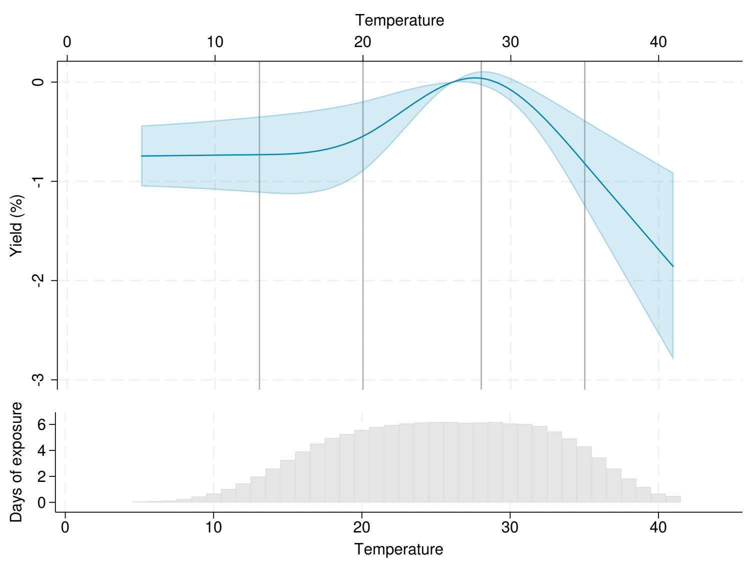

(b) Restricted spline with linear year trend

Figure A.5: Yield, robustness checks

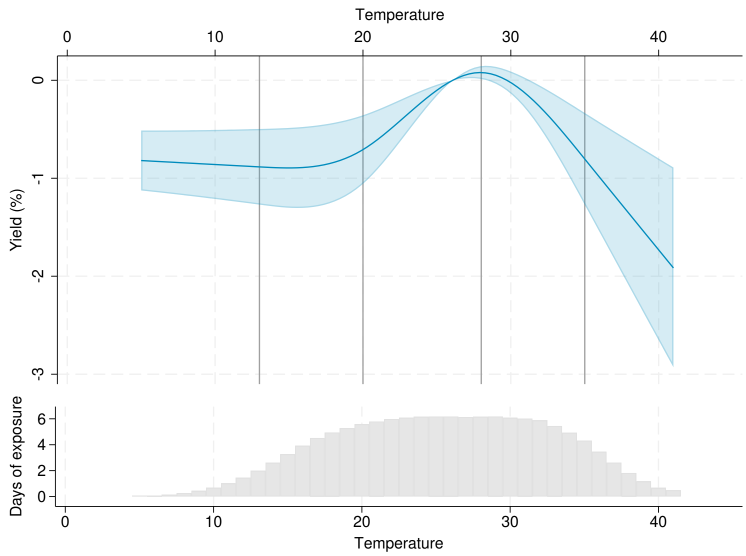

(c) Restricted spline with county-specific quadratic year trends

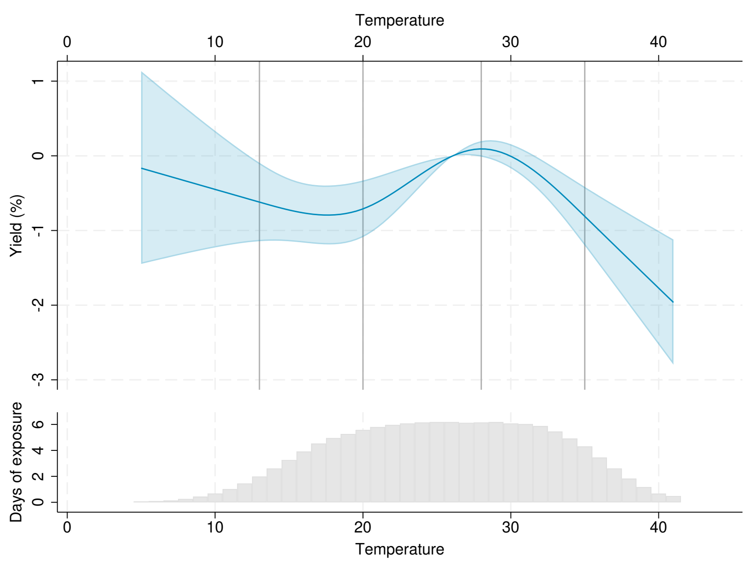

(d) Restricted spline with season length control

Figure A.5: Yield, robustness checks

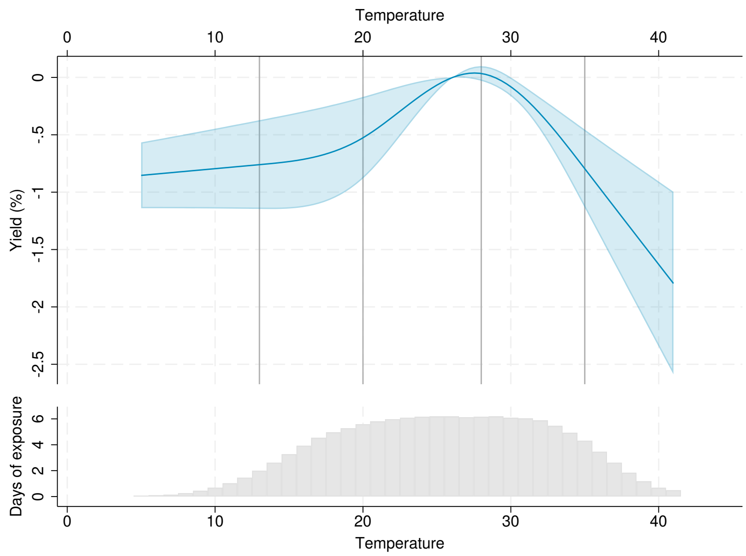

(e) Restricted spline without field-year controls

Figure A.5: Yield, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

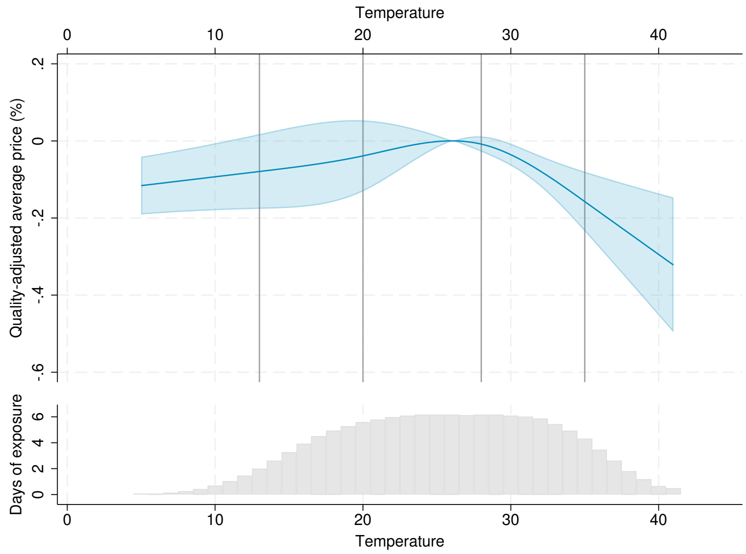

(a) Restricted spline with field fixed effects

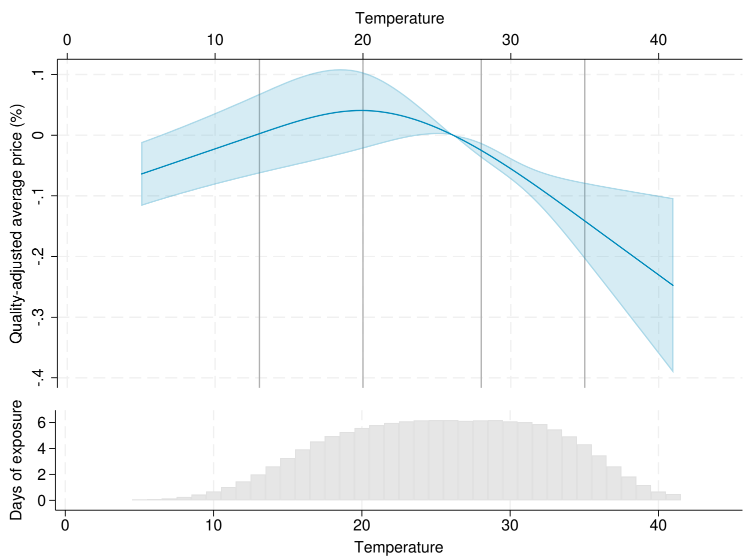

(b) Restricted spline with linear year trend

Figure A.6: Quality, robustness checks

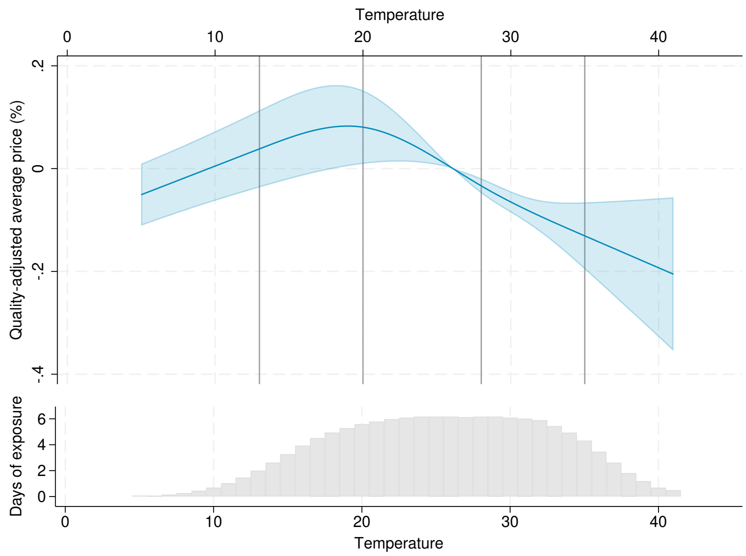

(c) Restricted spline with county-specific quadratic year trends

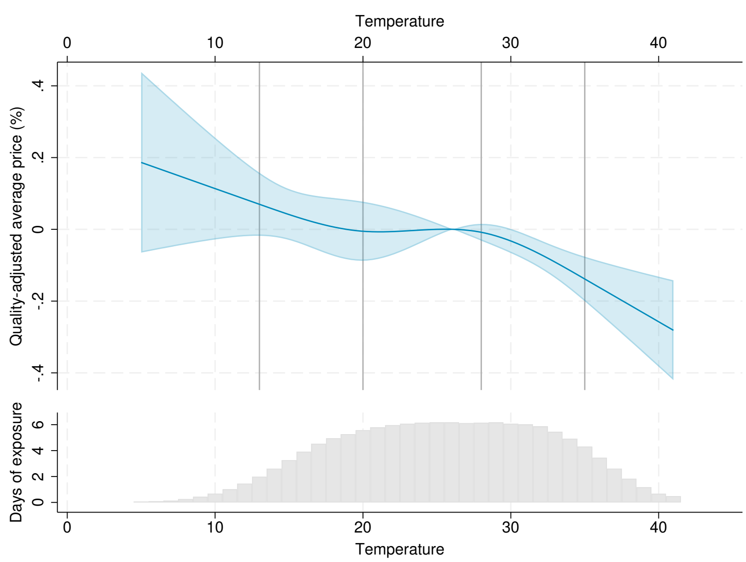

(d) Restricted spline with season length control

Figure A.6: Quality, robustness checks

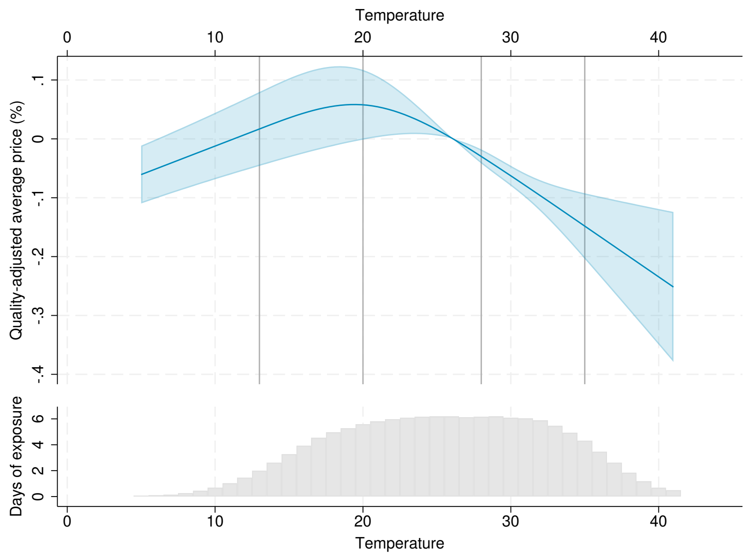

(e) Restricted spline without field-year controls

Figure A.6: Quality, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

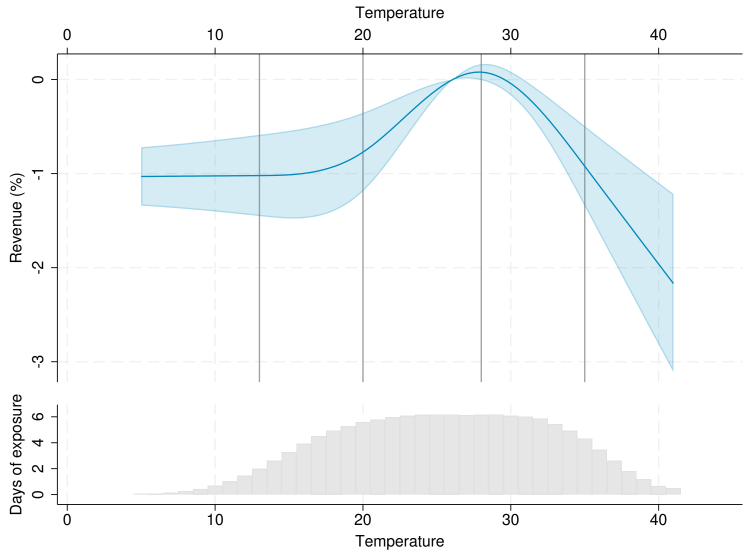

(a) Restricted spline with field fixed effects

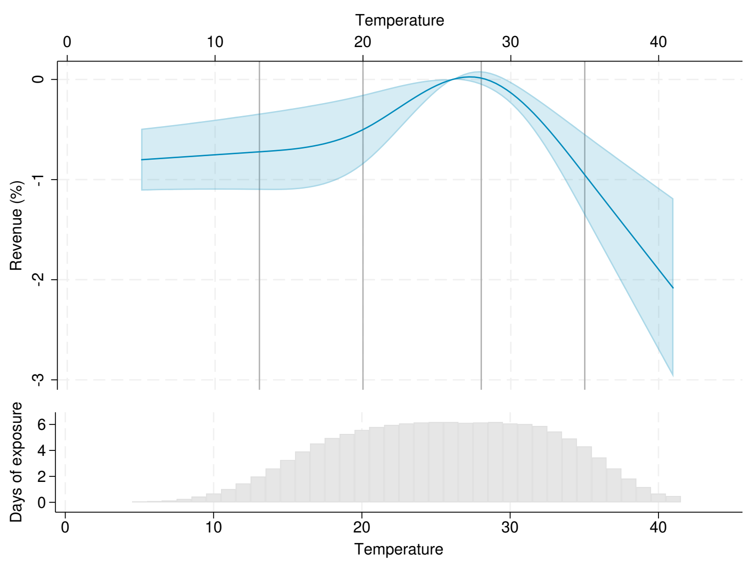

(b) Restricted spline with linear year trend

Figure A.7: Revenue per acre, robustness checks

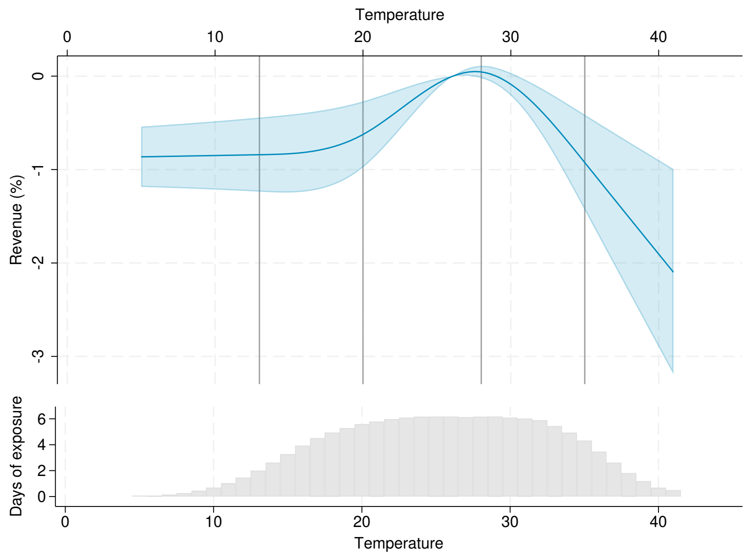

(c) Restricted spline with county-specific quadratic year trends

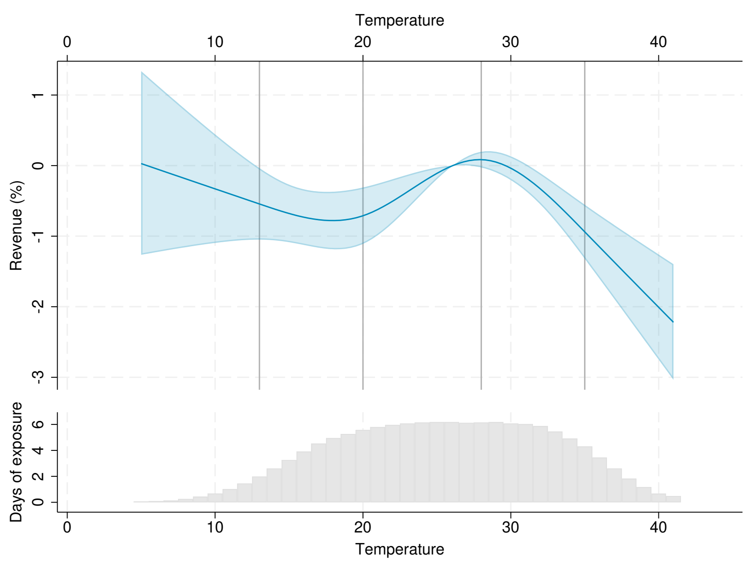

(d) Restricted spline with season length control

Figure A.7: Revenue per acre, robustness checks

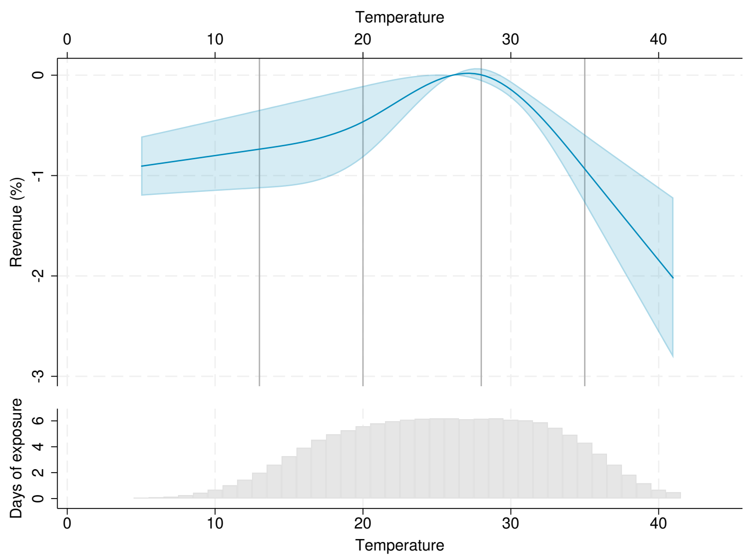

(e) Restricted spline without field-year controls

Figure A.7: Revenue per acre, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

D Variation in Yield, Quality, Revenue, and Weather

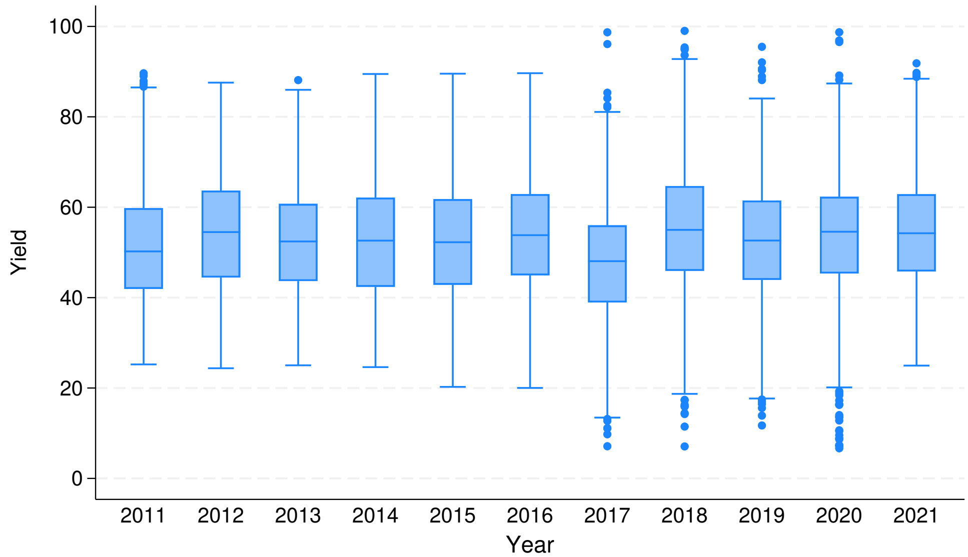

(a) Yield

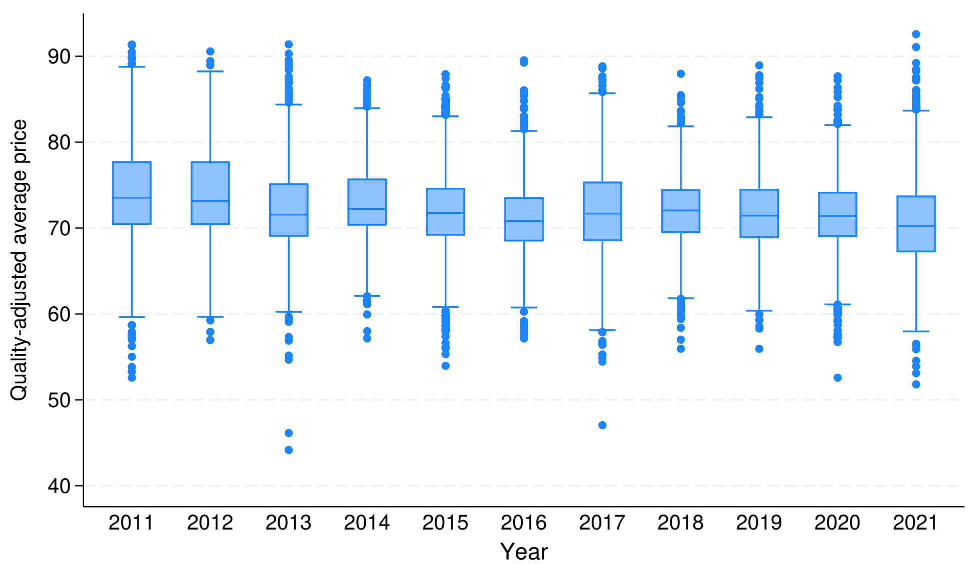

(b) Quality

Figure A.8: Distribution of yield, quality and revenue by year

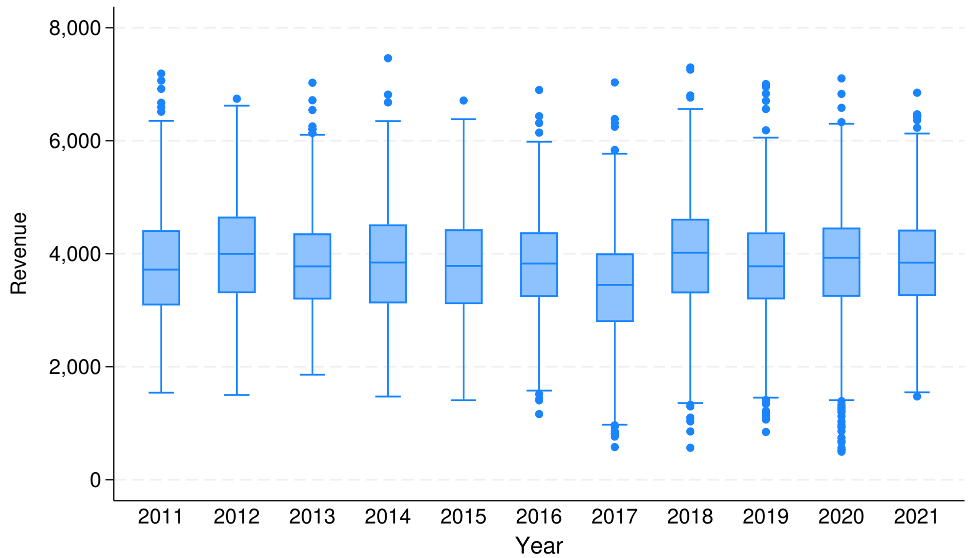

(c) Revenue

Figure A.8: Distribution of yield, quality and revenue by year

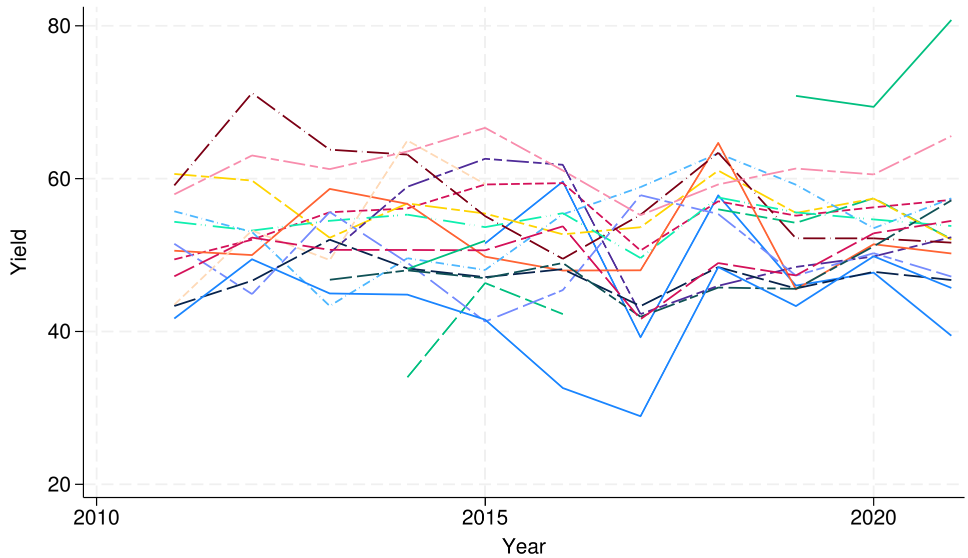

(d) Yield

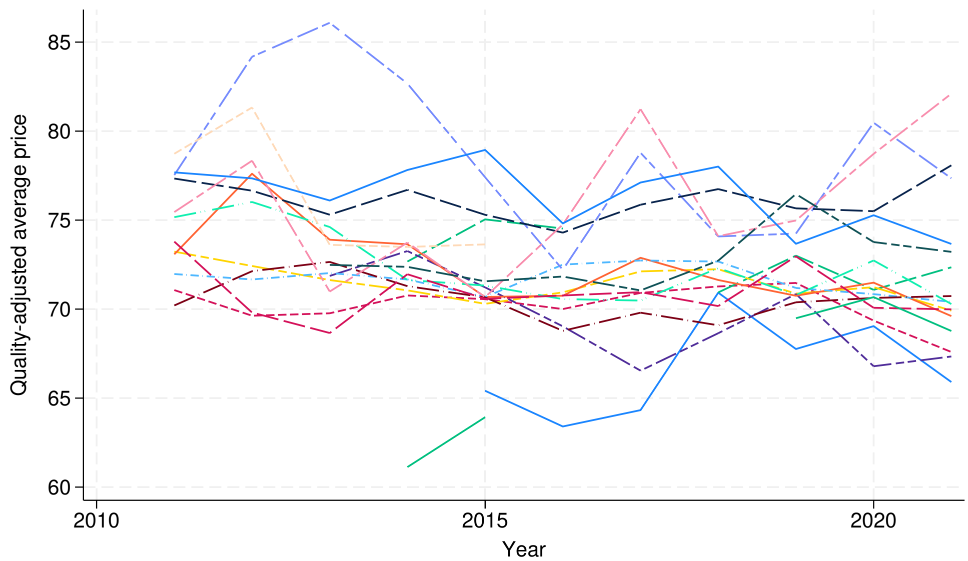

(e) Quality

Figure A.9: Average of yield, quality and revenue by year and county

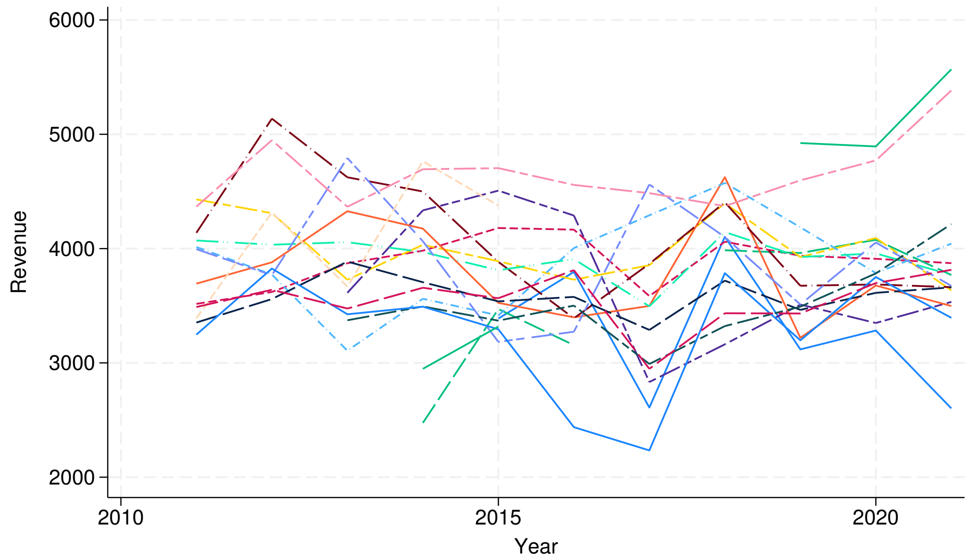

(f) Revenue

Figure A.9: Average of yield, quality and revenue by year and county

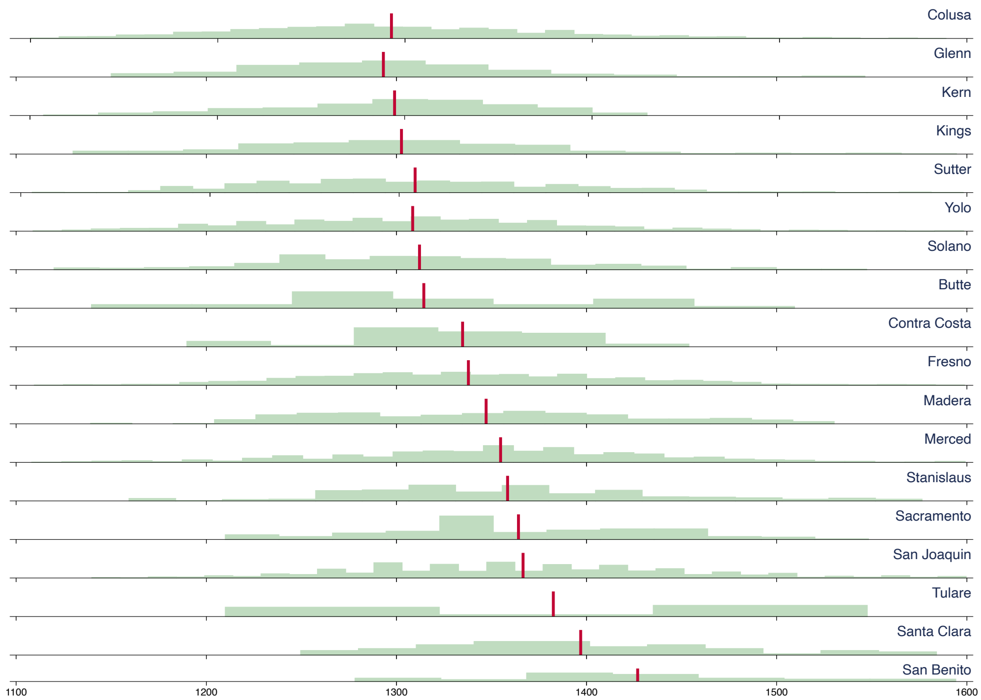

(g) Degree days below 10°C

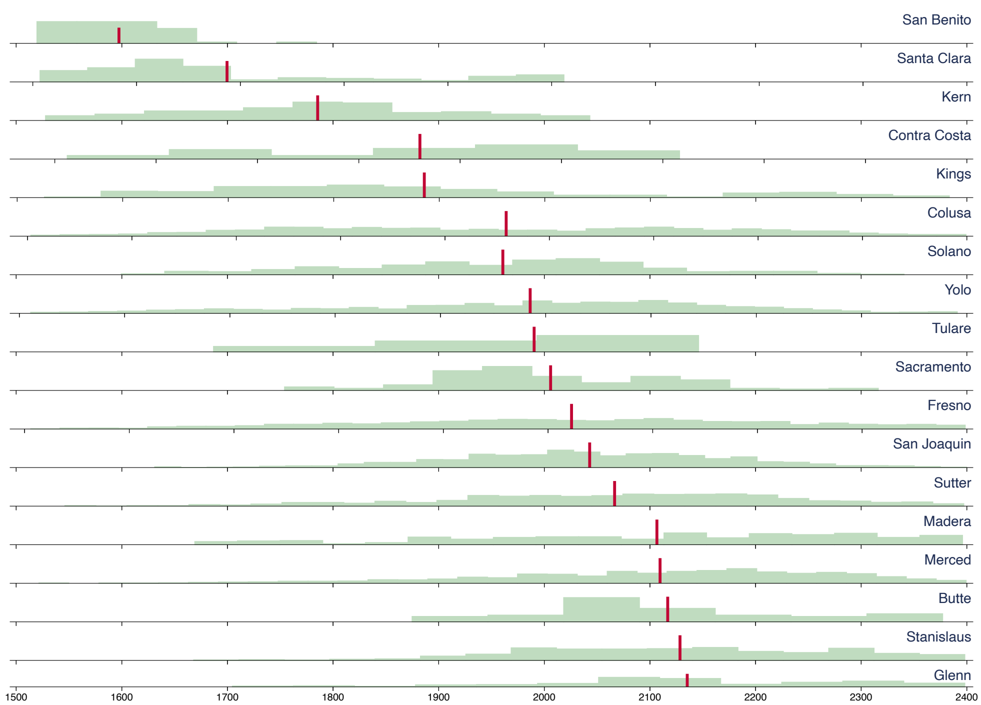

(h) Degree days between 10°C and 35°C

Figure A.10: Histogram of degree day variables by county, with county-average in red

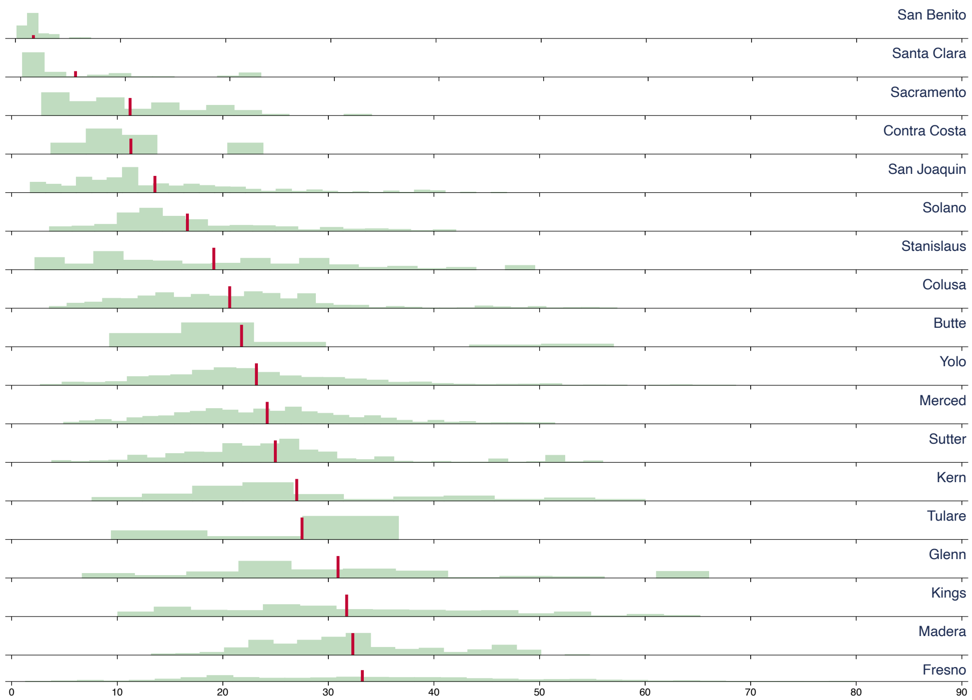

(i) Degree days above 35°C

Figure A.10: Histogram of degree day variables by county, with county-average in red

E Results Using Contemporaneous Prices

Our preferred specification uses a measure of price that removes variation in price (and revenue) that stems from market conditions and preserves variation caused by quality only. This avoids a potential issue of simultaneity bias. Here, we show results that use contemporaneous prices that incorporate the actual base price. Compared to other settings, simultaneity should not be as big of an issue in our setting because the annual base price is set in contracts before output for that year is realized. Results using contemporaneous prices are similar in direction but larger in magnitude than results from our preferred specification. This indicates that our preferred measure of price (and revenue) avoids some bias by removing variation from market conditions.

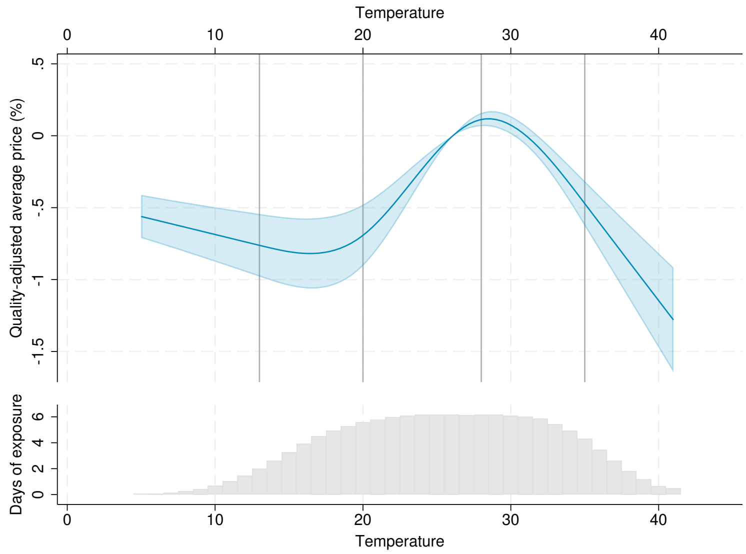

(j) Quality

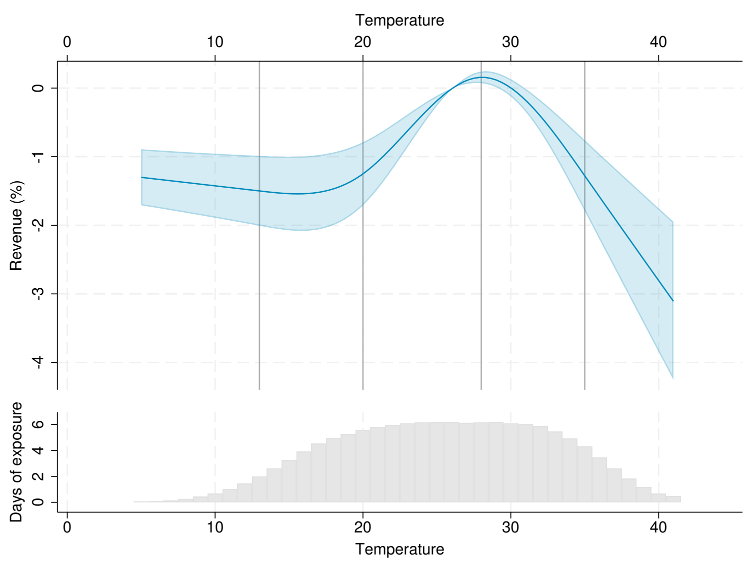

(k) Revenue per acre

Figure A.11: Results using contemporaneous prices

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

F Piecewise-Linear Degree-Day Specification

The piecewise-linear degree-day model is widely used in the agronomic and agricultural economic literature. It imposes more structure than flexible, semiparametric models. It also relies on the econometrician to correctly choose knot locations where the marginal effects change. However, it is less likely to overfit the data and has been shown in some contexts to perform better out-of-sample (Schlenker & Roberts, 2009).

To implement the piecewise-linear degree-day functional form, we first need to calculate degree days. Degree days are related to, but different from, temperature exposure. Temperature exposure measures how long is spent in a given temperature interval. Degree days measure how long and by how much temperatures exceed the lower bound of a temperature interval while being truncated at an upper bound (Snyder, 1985). When the temperature interval is small (e.g. 1°C), the difference between the two methods is relatively small because the "how much" dimension is unimportant relative to the "how long" dimension. When the temperature interval is large, as is the case in a piecewise-linear model, the difference between the two methods will be large. For example, if we use temperature exposure, we assume that the damage of one day of exposure at 35°C is equal to the damage of one day at 40°C. If we use degree days, we assume the damage of five days at 35°C is equal to the damage of one day at 40°C. The underlying assumption of degree days is that the effect of temperature exposure increases linearly with temperature between the lower and upper bounds.

Degree days can be computed from the temperature exposure vector x_{it}. The expression for calculating degree days between a lower bound of $\underline{h}$ and upper bound of $\overline{h}$ is:

$$DD_{it,[\underline{h},\overline{h}]} = \sum_{j=\underline{h}}^{\overline{h}-1} x_{it,j} \times (j - \underline{h} + 1) \tag{11}$$

Next, we choose knot locations. In the first set of results, we use knot locations suggested by the agronomic literature. Mid-range temperatures are ideal for yield and quality outcomes, but these outcomes may be damaged by hot (greater than 35°C) or cool (less than 10°C) temperatures (Hartz et al., 2008). Accordingly, we choose two knots at $\kappa_1 = 10$°C and $\kappa_2 = 35$°C. We estimate degree days using Equation 11 for each of the three "segments": below 10°C, between 10°C and 35°C, and above 35°C.

Equation 3 can then be modelled as:

$$y_{it} = \beta_0 + \beta_1 DD_{it,(-\infty,10]} + \beta_2 DD_{it,[10,35]} + \beta_3 DD_{it,[35,\infty)} + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{12}$$

In the second set of results, we use knot locations that correspond to the turning points in the restricted cubic-spline estimates.

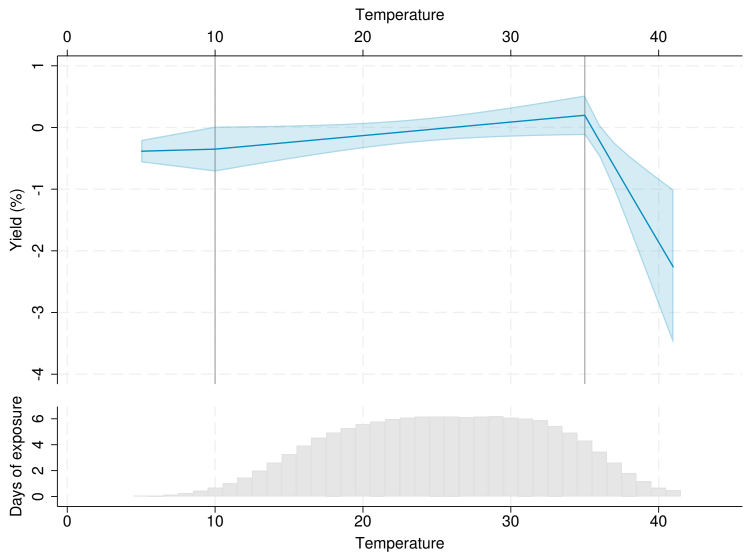

(a) Yield

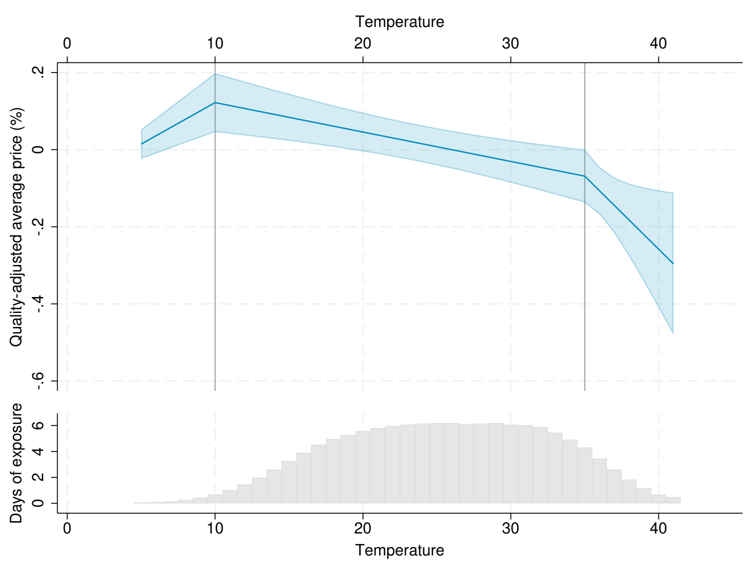

(b) Quality

Figure A.12: Piecewise-linear degree-day results, agronomic knots

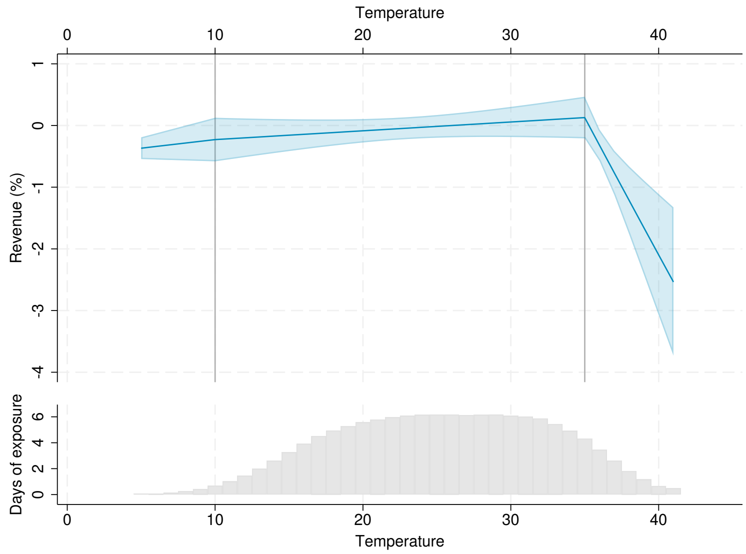

(c) Revenue per acre

Figure A.12: Piecewise-linear degree-day results, agronomic knots

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

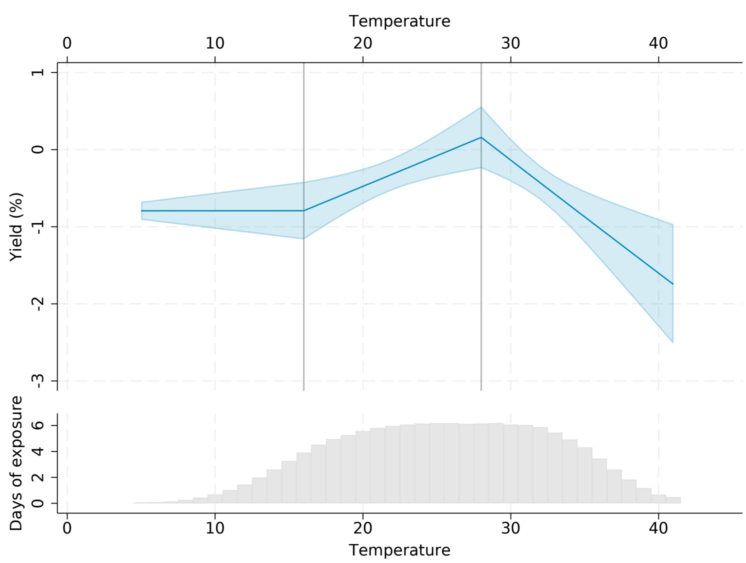

(a) Yield

Figure A.13: Piecewise-linear degree-day results, spline knots

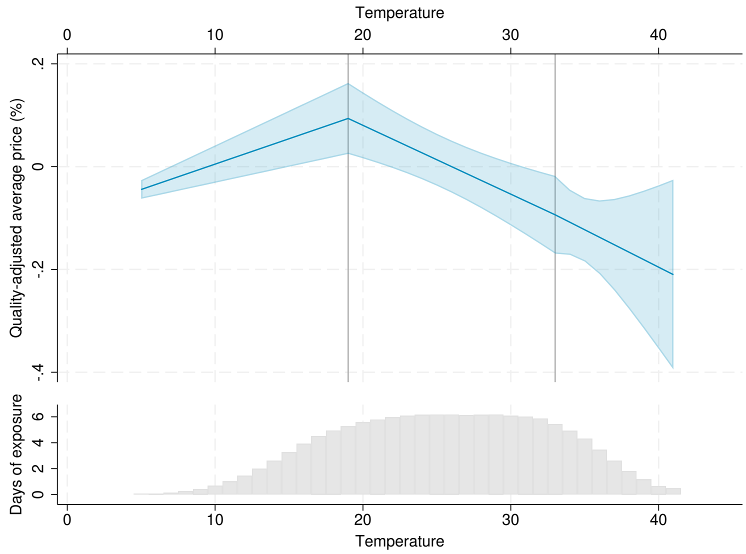

(b) Quality

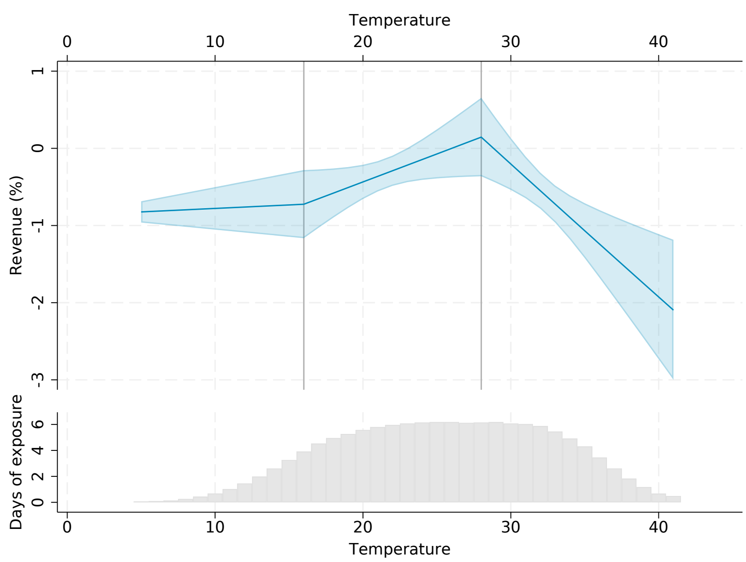

(c) Revenue per acre

Figure A.13: Piecewise-linear degree-day, spline knots

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

G Robustness of results to number of knots and knot placement

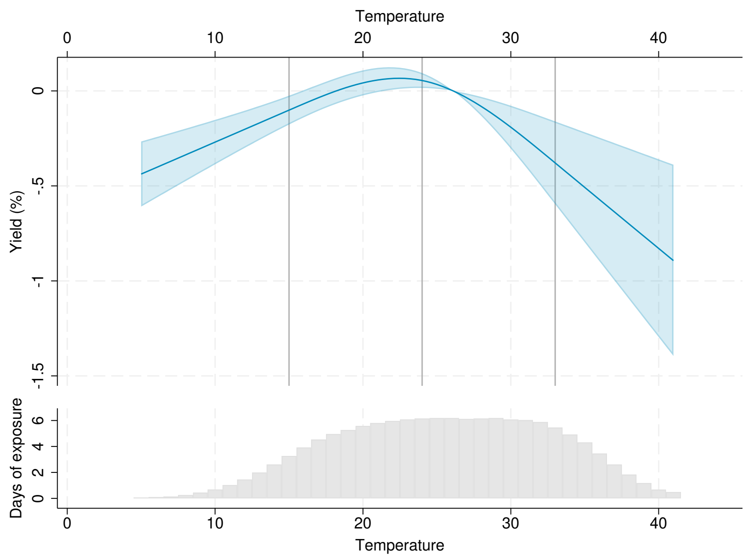

(d) Restricted spline with 3 knots

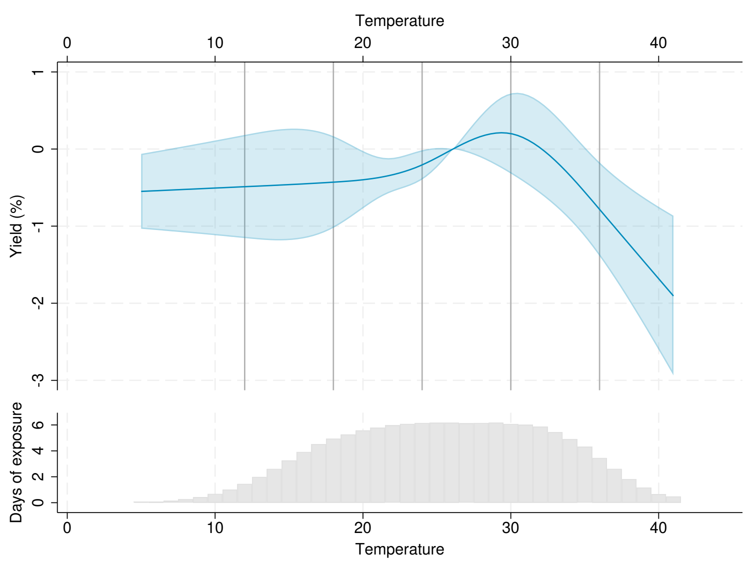

(e) Restricted spline with 5 knots

Figure A.14: Yield, robustness to varying knot number and placement

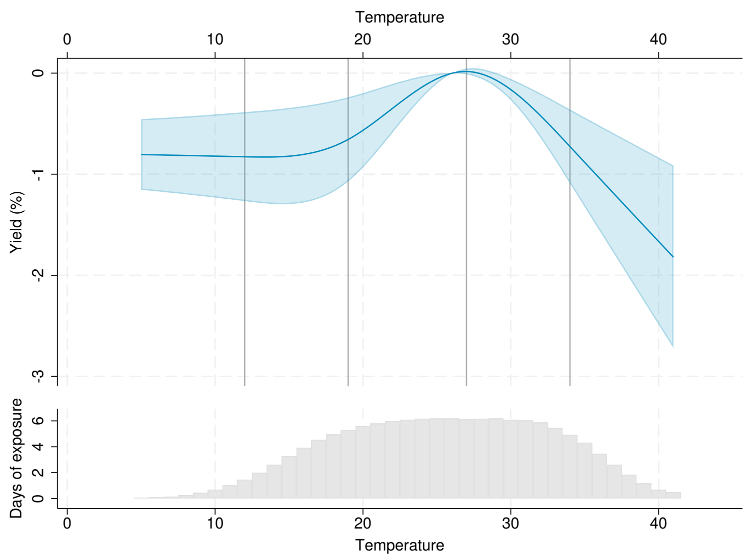

(f) Restricted spline with knots splitting the temperature series into quintiles

Figure A.14: Yield, robustness to varying knot number and placement

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

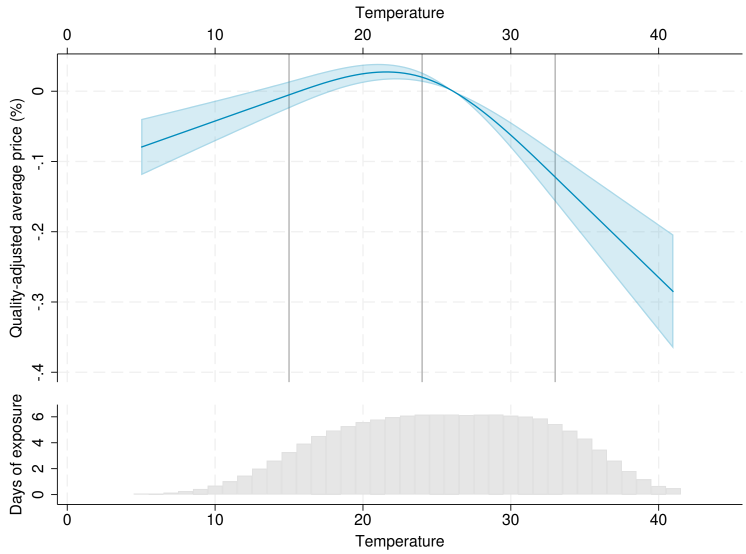

(a) Restricted spline with 3 knots

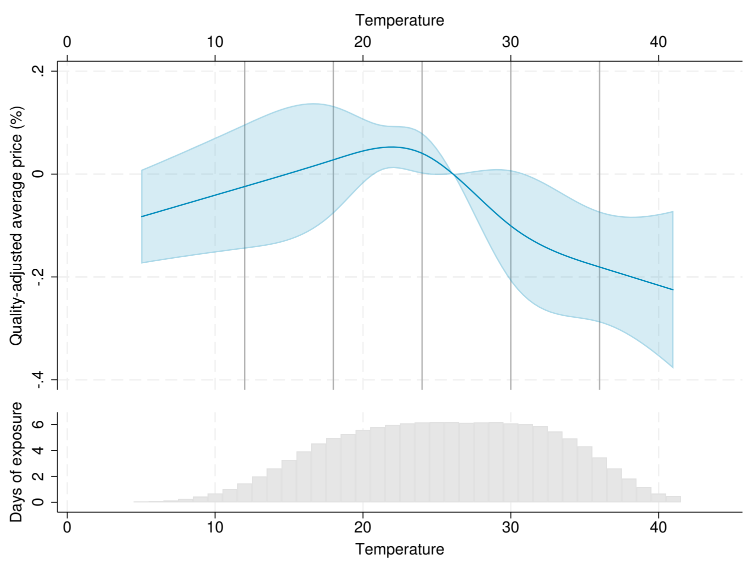

(b) Restricted spline with 5 knots

Figure A.15: Quality, robustness to varying knot number and placement

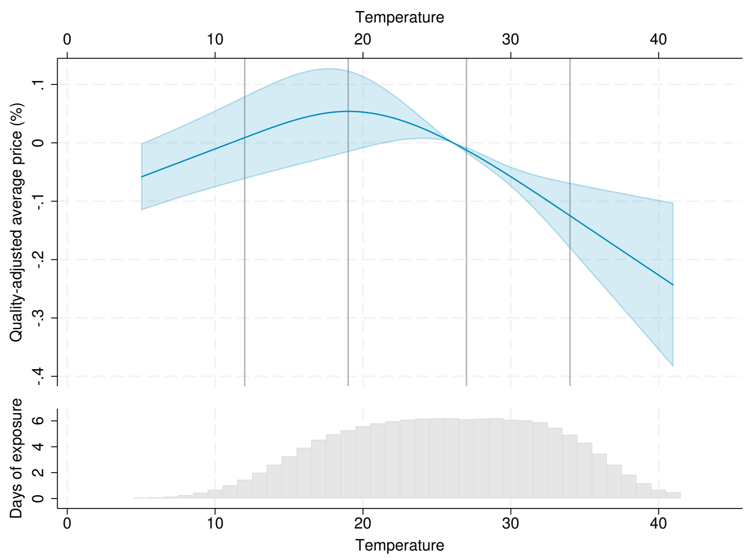

(c) Restricted spline with knots splitting the temperature series into quintiles

Figure A.15: Quality, robustness to varying knot number and placement

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

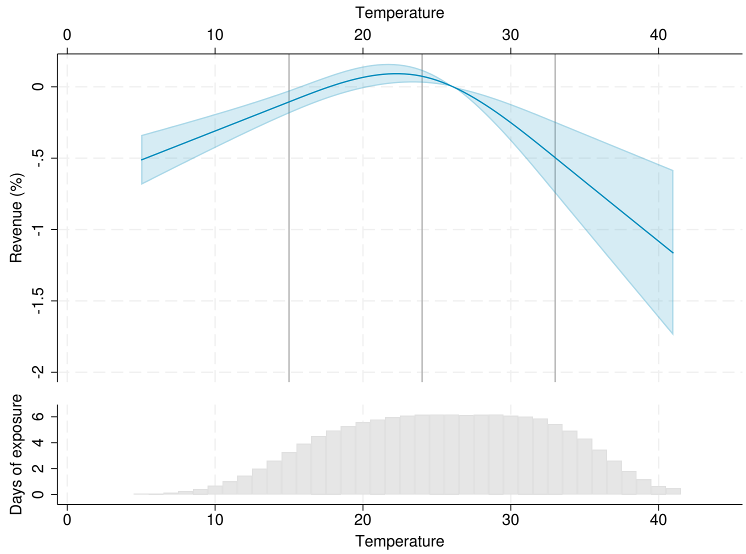

(a) Restricted spline with 3 knots

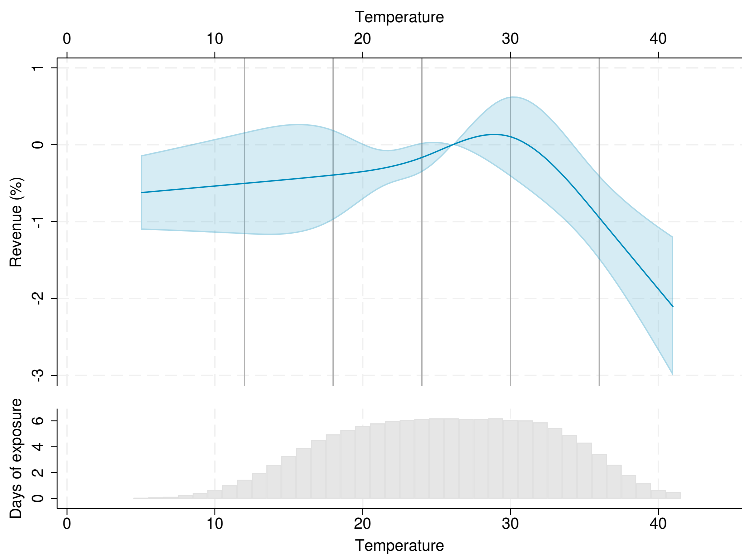

(b) Restricted spline with 5 knots

Figure A.16: Revenue per acre, robustness to varying knot number and placement

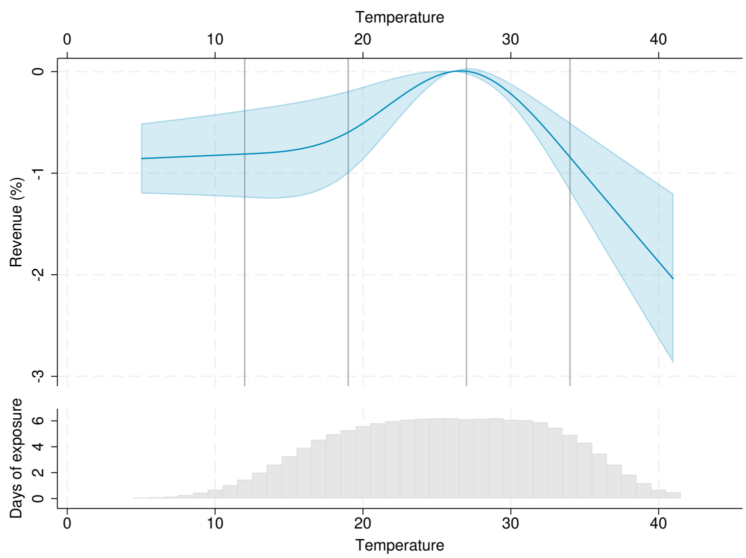

(c) Restricted spline with knots splitting the temperature series into quintiles

Figure A.16: Revenue per acre, robustness to varying knot number and placement

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

H Details on Restricted Cubic-Spline Specification

The B matrix for a cubic-spline combines the B matrix for the cubic polynomial and a $K \times J$ matrix Z. First, define $1 \times J$ temperature vector $W = \begin{pmatrix} 5 & 6 & 7 & \dots & 41 \end{pmatrix}$. Z is the resulting matrix after applying the cubic-spline basis function to W with K knots at $\kappa_1, \dots, \kappa_K$.

$$\begin{aligned} B_{cubic.spline} &= \begin{pmatrix} B_{cubic.polynomial} & Z \end{pmatrix} \\ &= \begin{pmatrix} 5 & 25 & 125 & (5-\kappa_1)^3_+ & (5-\kappa_2)^3_+ & \dots & (5-\kappa_K)^3_+ \\ 6 & 36 & 216 & (6-\kappa_1)^3_+ & (6-\kappa_2)^3_+ & \dots & (6-\kappa_K)^3_+ \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \\ 41 & 41^2 & 41^3 & (41-\kappa_1)^3_+ & (41-\kappa_2)^3_+ & \dots & (41-\kappa_K)^3_+ \end{pmatrix} \end{aligned} \tag{13}$$

Stone & Koo (1985) use the linearity constraints to develop a restricted cubic-spline function. Using this function, a restricted cubic-spline with K knots requires the estimation of only $K-1$ parameters on temperature (as opposed to $K+3$ parameters for the cubic spline). Equation 14 is the restricted spline function that we apply to W for $i = 1, 2, \dots, K-2$ [8].

$$\begin{aligned} V_1 &= W \\ V_{i+1} &= \frac{(W-\kappa_i)^3_+ - (\kappa_K - \kappa_{K-1})^{-1}\{(W-\kappa_{K-1})^3_+(\kappa_K-\kappa_i) - (W-\kappa_K)^3_+(\kappa_{K-1}-\kappa_i)\}}{(\kappa_K-\kappa_1)^2} \\ &\text{for } i = 1, 2, \dots, K-2 \end{aligned} \tag{14}$$

Together, W and V_2 to V_{K-1} make up the B matrix for the restricted cubic spline, where $V_{i,j}$ is the j-th element of the V_i vector.

$$\begin{aligned} B_{rest.cubic.spline} &= \begin{pmatrix} W & V_2 & \dots & V_{K-1} \end{pmatrix} \\ &= \begin{pmatrix} 5 & V_{2,0} & \dots & V_{K-1,0} \\ 6 & V_{2,1} & \dots & V_{K-1,1} \\ 7 & V_{2,2} & \dots & V_{K-1,2} \\ \vdots & \vdots & \ddots & \vdots \\ 41 & V_{2,45} & \dots & V_{K-1,45} \end{pmatrix} \end{aligned} \tag{15}$$

[8] For more details, see the Stata manual for the function mkspline

I Results for Individual Quality Attributes

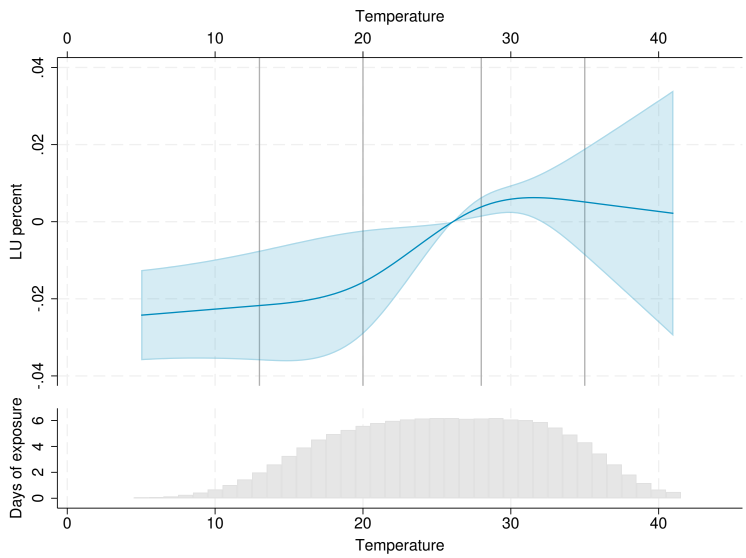

(a) Limited use (LU)

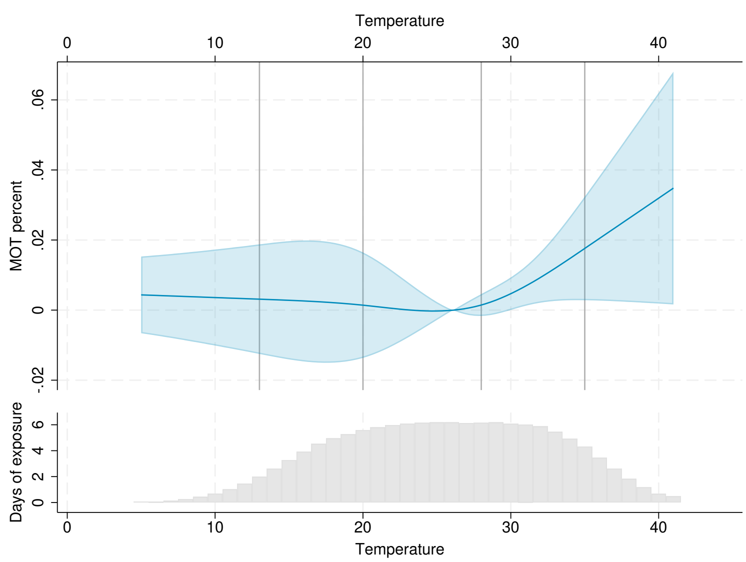

(b) Material other than tomatoes (MOT)

Figure A.17: Restricted cubic-spline results for individual quality defects

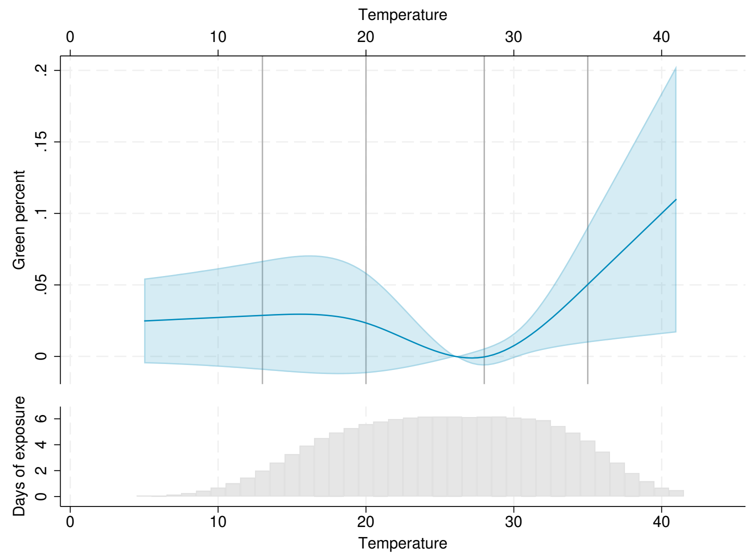

(c) Green

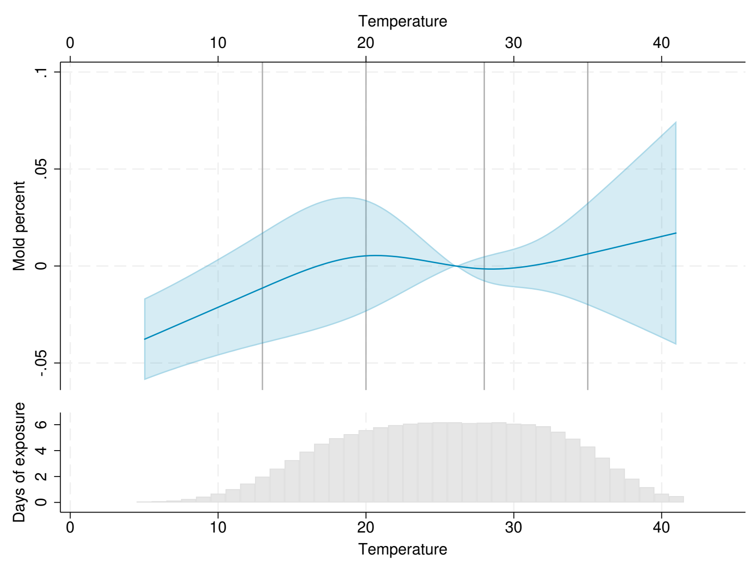

(d) Mold

Figure A.17: Restricted cubic-spline results for individual quality defects

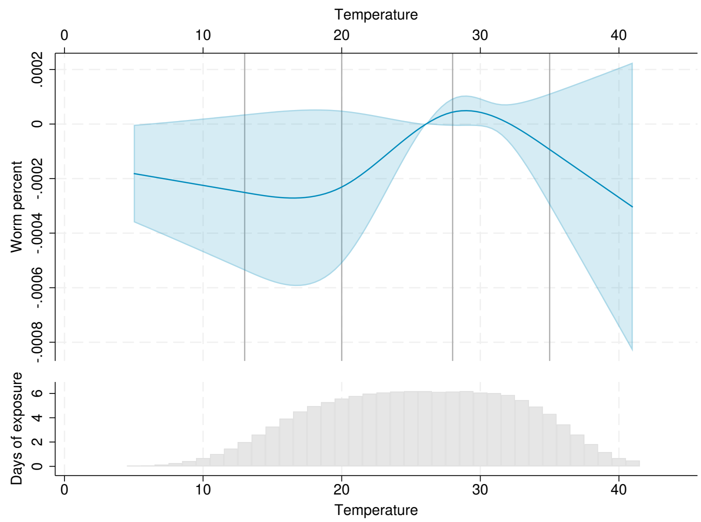

(e) Worm

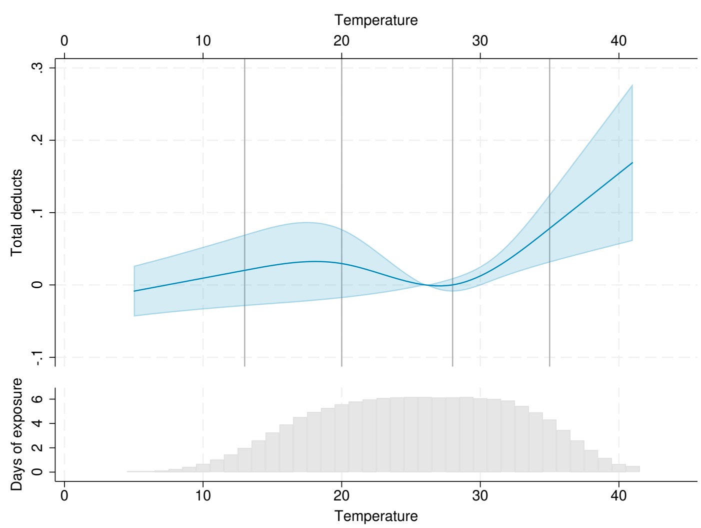

(f) Total defects

Figure A.17: Restricted cubic-spline results for individual quality defects

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

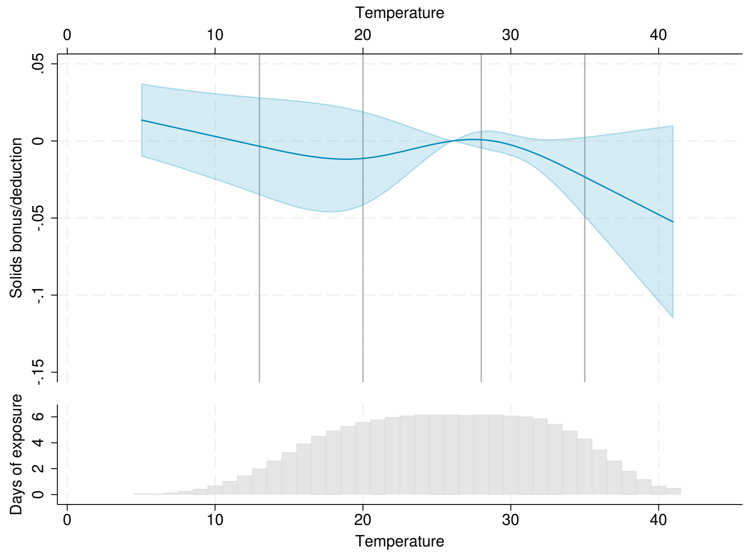

(a) Solids bonus

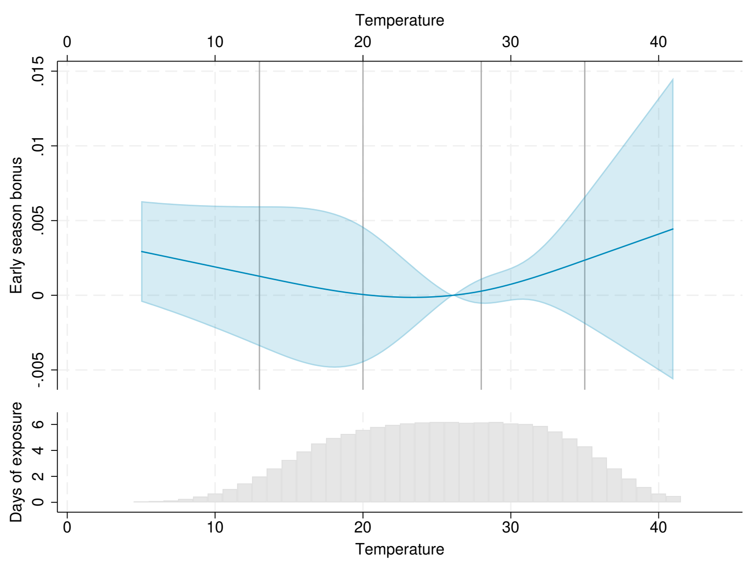

(b) Early season bonus

Figure A.18: Restricted cubic-spline results for individual quality bonuses

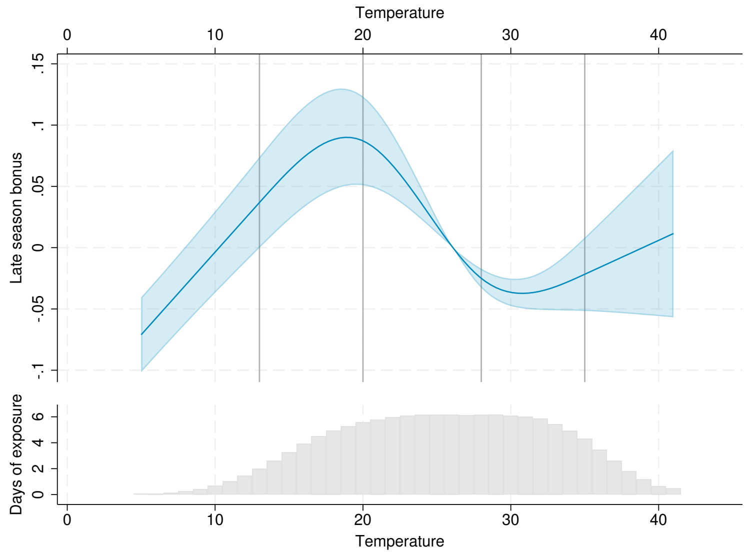

(c) Late season bonus

Figure A.18: Restricted cubic-spline results for individual quality bonuses

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

J Results for Quality Excluding Material Other than Tomatoes (MOT)

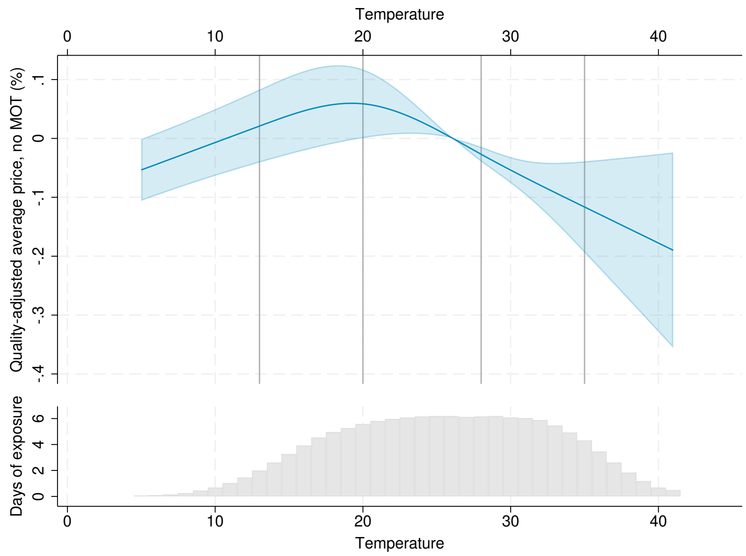

Figure A.19: Estimated effect of temperature on quality excluding MOT

Notes: The graph at the top of the frame shows the effect of an additional 24 hours in a given temperature interval on the outcome variable relative to 24 hours at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

K Estimated Effects of Control Variables

Table A1: Coefficient estimates on control variables

	(1) log yield	(2) log quality	(3) log revenue
Precipitation	-0.0009**	-0.0000	-0.0009**
	(0.0003)	(0.0001)	(0.0003)
Eolian	0.0495	-0.0034	0.0457
	(0.1177)	(0.0063)	(0.1134)
Organic material	-0.0275	-0.0057	-0.0330*
	(0.0152)	(0.0030)	(0.0158)
Lacustrine	-0.0066	0.0465***	0.0396
	(0.1455)	(0.0056)	(0.1410)
Residuum	0.0061	-0.0026	0.0035
	(0.0606)	(0.0049)	(0.0598)
Extended field storage variety	-0.0104	0.0188***	0.0082
	(0.0090)	(0.0026)	(0.0100)
Tomato spotted wilt resistant	0.0013	-0.0018	-0.0001
	(0.0074)	(0.0024)	(0.0082)
High solids	0.0262	-0.0233	0.0040
	(0.0482)	(0.0137)	(0.0498)
Fusarium Wilt resistant	-0.0297*	-0.0050	-0.0344*
	(0.0133)	(0.0034)	(0.0139)
Powdery Mildew resistant	-0.0456*	0.0135**	-0.0329
	(0.0193)	(0.0044)	(0.0203)
High yield	0.0118	0.0289*	0.0392
	(0.0496)	(0.0141)	(0.0520)
Fusarium Crown Rot resistant	0.0268	0.0149*	0.0418
	(0.0264)	(0.0059)	(0.0279)
Bacterial Spot resistant	-0.1334*	0.0101	-0.1236*
	(0.0520)	(0.0117)	(0.0549)
Early	-0.0479**	0.0195***	-0.0288
	(0.0153)	(0.0032)	(0.0153)
Thick	-0.0190	0.0027	-0.0167
	(0.0106)	(0.0025)	(0.0112)
Thin	-0.0539***	0.0048	-0.0490***
	(0.0145)	(0.0031)	(0.0141)
Pear-shaped	-0.0832	0.0023	-0.0810
	(0.0439)	(0.0060)	(0.0434)
Furrow irrigation	-0.1058***	0.0121*	-0.0938***
	(0.0170)	(0.0051)	(0.0171)
Sprinkler irrigation	-0.1593***	0.0128	-0.1457***
	(0.0101)	(0.0095)	(0.0170)
Harvesting early	-0.0021*	0.0001	-0.0020*
	(0.0009)	(0.0002)	(0.0008)
Year trend	-0.0061	0.0047	-0.0013
	(0.0116)	(0.0029)	(0.0123)
Year trend sqrd	0.0005	-0.0003	0.0001
	(0.0009)	(0.0002)	(0.0009)

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header. Standard errors (in parentheses) are two-way clustered by grower-group and county-year. Significance: * $p<0.01$, ** $p<0.05$, *** $p<0.001$.

L Statistical Significance of Selection Bias

To calculate whether selection introduces statistically significant bias in the estimated effect of temperature exposure on the outcomes, we estimate a model that nests our main specification Equation 5. We first append datasets with and without selection, and create a dummy variable equal to 1 for observations with added selection. For each outcome, we estimate a regression model that interacts all of the explanatory variables, control variables, and fixed effects with the selection dummy variable shown in Equation 16. We use our preferred method for calculating standard errors: heteroskedastic robust and two-way clustered by grower-group and county-by-year.

$$\begin{aligned} y_{it} &= x_{it}B\Gamma + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} + \\ &+ x_{it}B\Gamma' \times D_{it} + \delta' z_{it} \times D_{it} + \alpha'_{g(i)} \times D_{it} + \psi'(t) \times D_{it} + \epsilon'_{it} \times D_{it} \end{aligned} \tag{16}$$

Here, D_{it} is a dummy variable equal to 1 if the observation has added selection and 0 if the observation is part of the original dataset.

Finally, we do a F-test to test the hypothesis that the coefficients on splined temperature interacted with the selection dummy are jointly equal to zero $H_0 : \Gamma' = 0$. For yield and price, we reject the null hypothesis that the coefficients are jointly equal to zero and conclude that temperature has a differential effect on the outcomes in the presence of selection as shown in Table A2. For revenue, we fail to reject the null hypothesis and conclude that there is insufficient evidence to suggest temperature has a differential effect on revenue in the presence of selection—a consequence of the bias in yield and price canceling each other out.

Table A2: Results of testing the joint significance of the coefficients on the temperature exposure variables interacted with the selection dummy

	Yield	Quality	Revenue
F statistic	3.96	4.23	0.5275
p-value	0.0092	0.0064	0.74

Notes: F-test to test the hypothesis $H_0 : \Gamma' = 0$, where Γ' is a vector of the coefficients on splined temperature interacted with the selection dummy.

M Projected Impacts of Climate Change on Revenue

Following the methods outlined in Section 6, we also estimate the impact of climate change on revenue absent additional adaptation. Under both emissions scenarios, revenue is projected to decline from its 2011–21 levels by midcentury and end of century as shown in Figure A.20. However, these revenue results rely on the strong assumption that price incentives and contract structure remain fixed. In reality, the contract structure will evolve over time, likely in response to climate change. The projected impact of climate change on revenue should therefore be interpreted with a degree of caution.

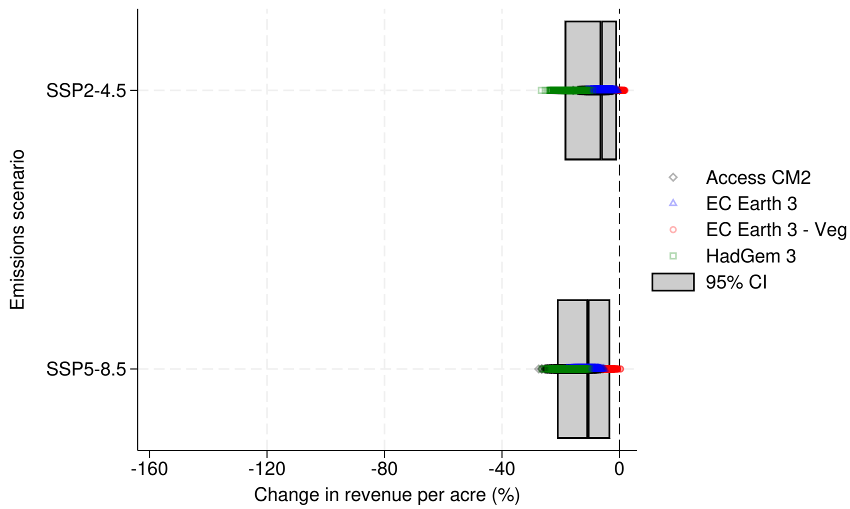

(a) Revenue per acre, mid-century

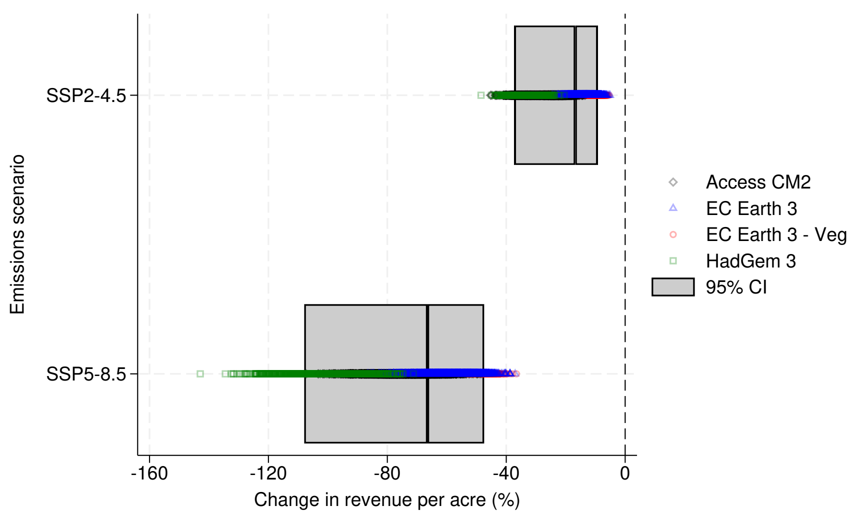

(b) Revenue per acre, end of century

Figure A.20: Projection of climate impacts by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

Notes: These graphs show the estimated impact of climate change on the outcome variables by midcentury and end of century, assuming no additional adaptation and all else equal. Each point is an estimate of the projected impact derived from a single combination of projection, emissions scenario, and wild-cluster bootstrap replication. The thick black lines represent the median impact estimate and the shaded grey areas represent the 95% confidence intervals that account for statistical and climate uncertainty.

N Projected Impacts of Climate Change with Selection on Quality

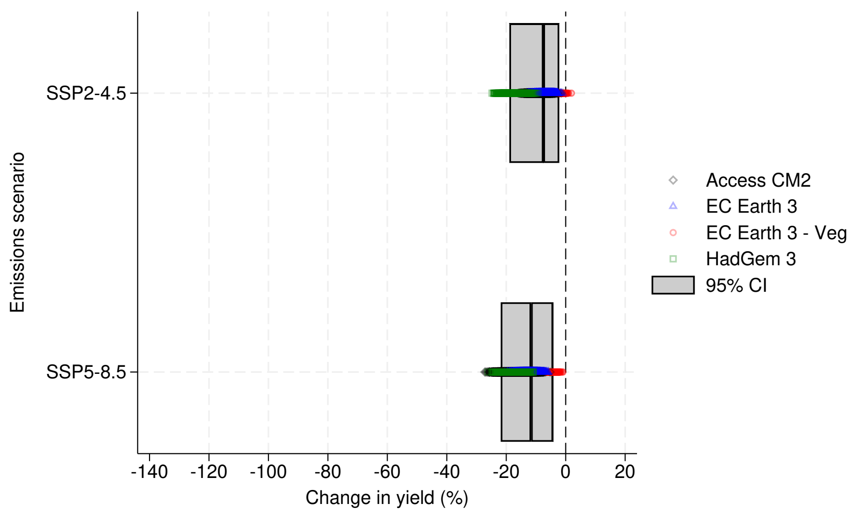

(a) Yield, mid-century

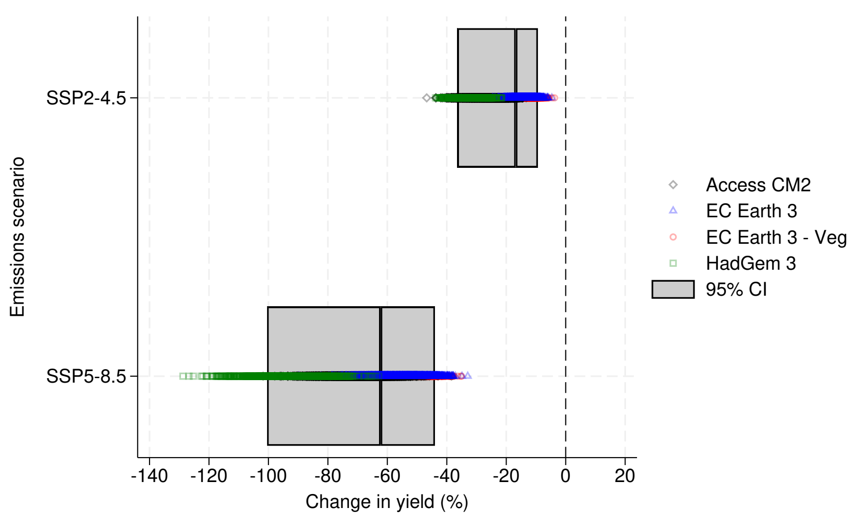

(b) Yield, end of century

Figure A.21: Projection of climate impacts with selection on quality by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

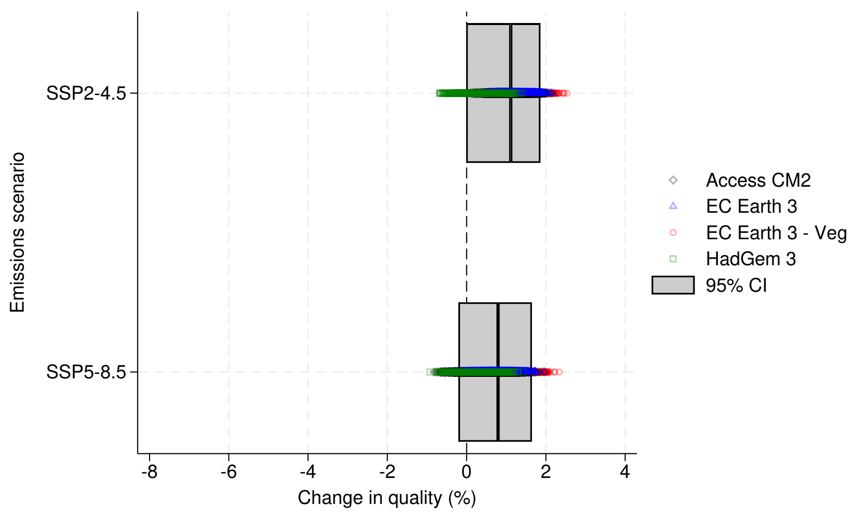

(c) Quality, mid-century

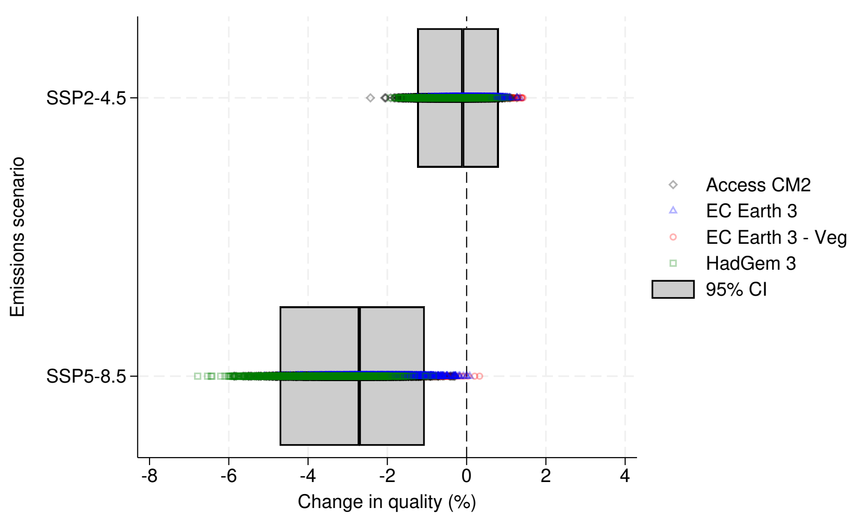

(d) Quality, end of century

Figure A.21: Projection of climate impacts with selection on quality by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

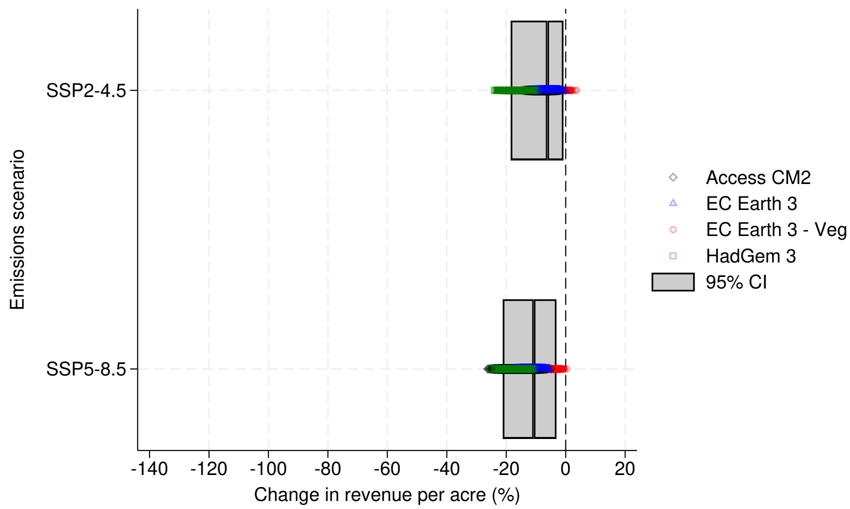

(e) Revenue per acre, mid-century

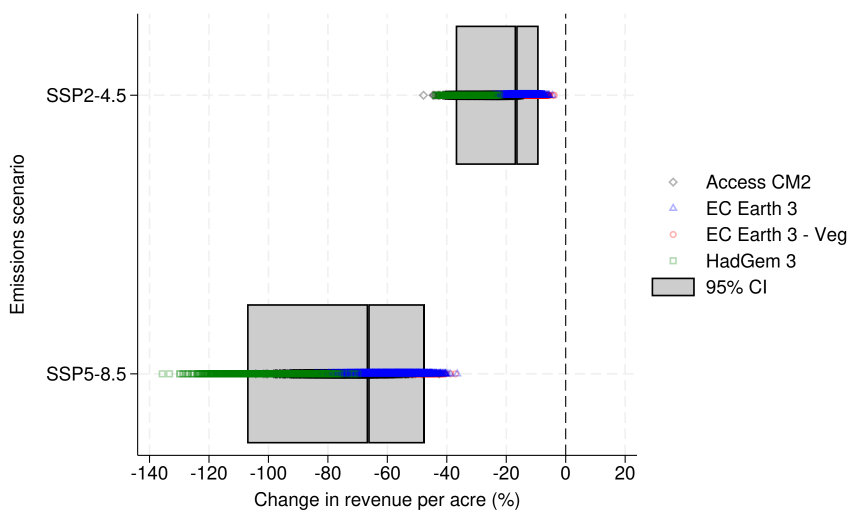

(f) Revenue per acre, end of century

Figure A.21: Projection of climate impacts with selection on quality by midcentury (2041–50) and end of century (2091–2100) relative to a 2011–21 baseline

Notes: These graphs show the estimated impact of climate change with selection on the outcome variables by midcentury and end of century, assuming no additional adaptation and all else equal. These results use yield and quality observations with added selection as experienced in other settings. Each point is an estimate of the projected impact derived from a single combination of projection, emissions scenario, and wild cluster bootstrap replication. The thick black lines represent the median impact estimate and the shaded grey areas represent the 95% confidence intervals that account for statistical and climate uncertainty.

References

AgSeeds (2020). Season Review. Available at `http://www.agseeds.com/`.

Ashenfelter, O. & Storchmann, K. (2016). Climate change and wine: A review of the economic implications. *Journal of Wine Economics*, 11(1), 105–138.

Bareille, F. & Chakir, R. (2023). The impact of climate change on agriculture: A repeat-Ricardian analysis. *Journal of Environmental Economics and Management*, 119, 102822.

Berry, S. T., Roberts, M. J., & Schlenker, W. (2014). Corn production shocks in 2012 and beyond: Implications for harvest volatility. In *The Economics of Food Price Volatility* (pp. 59–81). University of Chicago Press.

Blanc, E. & Schlenker, W. (2020). The use of panel models in assessments of climate impacts on agriculture. *Review of Environmental Economics and Policy*.

Bucheli, J., Dalhaus, T., & Finger, R. (2022). Temperature effects on crop yields in heat index insurance. *Food Policy*, 107, 102214.

Burke, M., Dykema, J., Lobell, D. B., Miguel, E., & Satyanath, S. (2015). Incorporating climate uncertainty into estimates of climate change impacts. *Review of Economics and Statistics*, 97(2), 461–471.

Cal-Adapt (2023). Data derived from LOCA2 Downscaled CMIP6 Climate Projections. Cal-Adapt website developed by University of California at Berkeley's Geospatial Innovation Facility under contract with the California Energy Commission. Available at `https://cal-adapt.org/`.

California Department of Food and Agriculture (1997). California Processing Tomato Inspection Program. Available at `https://www.cdfa.ca.gov/mkt/mkt/pdf/ProcessingTomatoInspection.pdf`.

California State Geoportal (2023). Water Districts. Available at `https://gis.data.ca.gov/datasets/45d26a15b96346f1816d8fe187f8570d_0/about`.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3), 414–427.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249.

Cammarano, D., Jamshidi, S., Hoogenboom, G., Ruane, A. C., Niyogi, D., & Ronga, D. (2022). Processing tomato production is expected to decrease by 2050 due to the projected increase in temperature. *Nature Food*, (pp. 1–8).

Carter, C., Cui, X., Ghanem, D., & Mérel, P. (2018). Identifying the economic impacts of climate change on agriculture. *Annual Review of Resource Economics*, 10, 361–380.

Carter, E. K., Melkonian, J., Riha, S. J., & Shaw, S. B. (2016). Separating heat stress from moisture stress: analyzing yield response to high temperature in irrigated maize. *Environmental Research Letters*, 11(9), 094012.

CDFA (2022). 2021-2022 California Agricultural Statistics Review. Available at `https://www.cdfa.ca.gov/Statistics/PDFs/2022_Ag_Stats_Review.pdf`.

Chen, S., Chen, X., & Xu, J. (2016). Impacts of climate change on agriculture: Evidence from china. *Journal of Environmental Economics and Management*, 76, 105–124.

Cheng, W., Dan, L., Deng, X., Feng, J., Wang, Y., Peng, J., Tian, J., Qi, W., Liu, Z., Zheng, X., et al. (2022). Global monthly gridded atmospheric carbon dioxide concentrations under the historical and future scenarios. *Scientific Data*, 9(1), 83.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of econometrics*, 92(1), 1–45.

Cui, X. (2020). Beyond yield response: weather shocks and crop abandonment. *Journal of the Association of Environmental and Resource Economists*, 7(5), 901–932.

Cui, X. & Xie, W. (2022). Adapting agriculture to climate change through growing season adjustments: Evidence from corn in china. *American Journal of Agricultural Economics*, 104(1), 249–272.

D'Agostino, A. L. & Schlenker, W. (2016). Recent weather fluctuations and agricultural yields: implications for climate change. *Agricultural economics*, 47(S1), 159–171.

Dalhaus, T., Schlenker, W., Blanke, M. M., Bravin, E., & Finger, R. (2020). The effects of extreme weather on apple quality. *Scientific reports*, 10(1), 1–7.

Deschênes, O. & Greenstone, M. (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American Economic Review*, 97(1), 354–385.

Elliott, J., Deryng, D., Müller, C., Frieler, K., Konzmann, M., Gerten, D., Glotter, M., Flörke, M., Wada, Y., Best, N., et al. (2014). Constraints and potentials of future irrigation water availability on agricultural production under climate change. *Proceedings of the National Academy of Sciences*, 111(9), 3239–3244.

FAO (2023). Food and Agriculture Organization Corporate Statistical Database (FAOSTAT). Available at `https://www.fao.org/faostat/en/`.

Fezzi, C. & Bateman, I. (2015). The impact of climate change on agriculture: nonlinear effects and aggregation bias in Ricardian models of farmland values. *Journal of the Association of Environmental and Resource Economists*, 2(1), 57–92.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., & Schlenker, W. (2012). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *American Economic Review*, 102(7), 3749–3760.

Gammans, M., Mérel, P., & Ortiz-Bobea, A. (2017). Negative impacts of climate change on cereal yields: statistical evidence from France. *Environmental Research Letters*, 12(5), 054007.

Hartz, T., Miyao, G., Mickler, J., Lestrange, M., Stoddard, S., Nuñez, J., & Aegerter, B. (2008). Processing tomato production in California. *UC Vegetable Research & Information Center*.

Hayhoe, K., Cayan, D., Field, C. B., Frumhoff, P. C., Maurer, E. P., Miller, N. L., Moser, S. C., Schneider, S. H., Cahill, K. N., Cleland, E. E., et al. (2004). Emissions pathways, climate change, and impacts on California. *Proceedings of the national academy of sciences*, 101(34), 12422–12427.

Hendricks, N. P. (2018). Potential benefits from innovations to reduce heat and water stress in agriculture. *Journal of the Association of Environmental and Resource Economists*, 5(3), 545–576.

Hill, A. E. & Beatty, T. K. (2024). Evidence on quality spillovers from speed enhancing policies in the workplace. *Economic Inquiry*, 62(4), 1520–1538.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of sciences*, 107(35), 15367–15372.

IPCC (2023). *Sixth Assessment (AR6) Synthesis Report: Climate Change 2023*. Cambridge University Press.

Just, R. E. & Chern, W. S. (1980). Tomatoes, technology, and oligopsony. *The Bell Journal of Economics*, (pp. 584–602).

Kawasaki, K. (2019). Two harvests are better than one: double cropping as a strategy for climate change adaptation. *American Journal of Agricultural Economics*, 101(1), 172–192.

Kawasaki, K. (2023). Impact of climate change on crop pests and diseases: Ensemble modeling of time-varying weather effects. *Journal of the Association of Environmental and Resource Economists*, 10(6), 1515–1543.

Kawasaki, K. & Lichtenberg, E. (2014). Econometric analysis of grading standards: The ordered fractional approach. *American Journal of Agricultural Economics*, 96(1), 345–365.

Kawasaki, K. & Uchida, S. (2016). Quality matters more than quantity: asymmetric temperature effects on crop yield and quality grade. *American Journal of Agricultural Economics*, 98(4), 1195–1209.

Krantz, W., Pierce, D., Goldenson, N., & Cayan, D. (2021). Memorandum on Evaluating Global Climate Models for Studying Regional Climate Change in California. *Scripps Institution of Oceanography, UC San Diego*.

Liang, X.-Z., Wu, Y., Chambers, R. G., Schmoldt, D. L., Gao, W., Liu, C., Liu, Y.-A., Sun, C., & Kennedy, J. A. (2017). Determining climate effects on US total agricultural productivity. *Proceedings of the National Academy of Sciences*, 114(12), E2285–E2292.

Lobell, D. B., Cahill, K. N., & Field, C. B. (2007). Historical effects of temperature and precipitation on California crop yields. *Climatic change*, 81(2), 187–203.

Lobell, D. B., Torney, A., & Field, C. B. (2011). Climate extremes in California agriculture. *Climatic change*, 109(1), 355–363.

Marklein, A., Elias, E., Nico, P., & Steenwerth, K. (2020). Projected temperature increases may require shifts in the growing season of cool-season crops and the growing locations of warm-season crops. *Science of The Total Environment*, 746, 140918.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1994). The impact of global warming on agriculture: a Ricardian analysis. *The American economic review*, (pp. 753–771).

Naylor, R. L., Battisti, D. S., Vimont, D. J., Falcon, W. P., & Burke, M. B. (2007). Assessing risks of climate variability and climate change for Indonesian rice agriculture. *Proceedings of the National Academy of Sciences*, 104(19), 7752–7757.

NRCS USDA (2020). Soil Survey Geographic (SSURGO) Database for California. Available at `https://websoilsurvey.nrcs.usda.gov/`.

Ortiz-Bobea, A. (2021). The empirical analysis of climate change impacts and adaptation in agriculture. In *Handbook of agricultural economics*, volume 5 (pp. 3981–4073). Elsevier.

Ortiz-Bobea, A., Knippenberg, E., & Chambers, R. G. (2018). Growing climatic sensitivity of US agriculture linked to technological change and regional specialization. *Science advances*, 4(12), eaat4343.

Ortiz-Bobea, A., Wang, H., Carrillo, C. M., & Ault, T. R. (2019). Unpacking the climatic drivers of US agricultural yields. *Environmental Research Letters*, 14(6), 064003.

Pierce, D., F., K. J., & Cayan, D. (2023). Localized Construction Analogs. Scripps Institution of Oceanography.

PRISM (2020). PRISM Climate Data. Available at `http://prism.oregonstate.edu`.

Ramsey, A. F. & Rejesus, R. M. (2021). Bayesian hierarchical models for measuring varietal improvement in tobacco yield and quality. *Journal of Agricultural and Applied Economics*, 53(4), 563–586.

Ramsey, A. F., Tack, J. B., & Balota, M. (2020). Double or nothing: Impacts of warming on crop quantity, quality, and revenue. *Journal of Agricultural and Resource Economics*, (pp. 1–22).

Rangaswamy, T. C., Sridhara, S., Ramesh, N., Gopakkali, P., El-Ansary, D. O., Mahmoud, E. A., Abdelmohsen, S. A., Abdelbacki, A. M., Elansary, H. O., & Abdel-Hamid, A. M. (2021). Assessing the impact of higher levels of co2 and temperature and their interactions on tomato (solanum lycopersicum l.). *Plants*, 10(2), 256.

Schlenker, W. & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. *Proceedings of the National Academy of sciences*, 106(37), 15594–15598.

Schmitt, J., Offermann, F., Söder, M., Frühauf, C., & Finger, R. (2022). Extreme weather events cause significant crop yield losses at the farm level in German agriculture. *Food Policy*, 112, 102359.

Shaw, S. B., Mehta, D., & Riha, S. J. (2014). Using simple data experiments to explore the influence of non-temperature controls on maize yields in the mid-West and Great Plains. *Climatic change*, 122(4), 747–755.

Shew, A. M., Tack, J. B., Nalley, L. L., & Chaminuka, P. (2020). Yield reduction under climate warming varies among wheat cultivars in South Africa. *Nature communications*, 11(1), 1–9.

Snyder, R. L. (1985). Hand calculating degree days. *Agricultural and forest meteorology*, 35(1-4), 353–358.

Stone, C. J. & Koo, C.-Y. (1985). Additive splines in statistics. In *Proceedings of the Statistical Computing Section ASA. Washington, DC* (pp. 45–48).

Tack, J., Barkley, A., & Hendricks, N. (2017a). Irrigation offsets wheat yield reductions from warming temperatures. *Environmental Research Letters*, 12(11), 114027.

Tack, J., Barkley, A., & Nalley, L. L. (2015). Effect of warming temperatures on US wheat yields. *Proceedings of the National Academy of Sciences*, 112(22), 6931–6936.

Tack, J., Lingenfelser, J., & Jagadish, S. K. (2017b). Disaggregating sorghum yield reductions under warming scenarios exposes narrow genetic diversity in US breeding programs. *Proceedings of the National Academy of Sciences*, 114(35), 9296–9301.

USDA ERS (2020). Ag and Food Statistics: Charting the Essentials, February 2020. Available at `https://www.ers.usda.gov/webdocs/publications/96957/ap-083.pdf?v=1384.6`.

USDA ERS & NASS (2023). Percent of farm grain production used on farm, 1996–2019. ERS tabulation of microdata from the Agricultural Resource Management Survey, 1996–2019.

USDA NASS (2017). 2017 Census of Agriculture. Available at `www.nass.usda.gov/AgCensus`.

USDA NASS (2021a). 2020 California Processing Tomato County Estimates. Available at `https://www.nass.usda.gov/Statistics_by_State/California/Publications/County_Estimates/2021/PTOMCounty_0321.pdf`.

USDA NASS (2021b). 2021 California Processing Tomato Report. Available at `https://www.nass.usda.gov/Statistics_by_State/California/Publications/Specialty_and_Other_Releases/Tomatoes/2021/PTOMInt_0121.pdf`.

Wing, I. S., De Cian, E., & Mistry, M. N. (2021). Global vulnerability of crop yields to climate change. *Journal of Environmental Economics and Management*, 109, 102462.

Wu, X., Yu, L., & Pehrsson, P. R. (2022). Are processed tomato products as nutritious as fresh tomatoes? Scoping review on the effects of industrial processing on nutrients and bioactive compounds in tomatoes. *Advances in Nutrition*, 13(1), 138–151.